

ONLINE ADVERTISING TECHNOLOGY FOR ARTIFICIAL GENERAL INTELLIGENCE
(AGI) AND SUPERINTELLIGENCE (SI)

1.0 TECHNICAL FIELD

5 **[0001]** In some aspects, the present technology relates to an online advertising technology for Artificial Intelligence (AI), Artificial General Intelligence (AGI) and SuperIntelligence AGI (SuperIntelligence or “SI”) for use in connection with increasing the intelligence of AI/AGI systems by overcoming the data bottleneck that AI researchers currently face.

10 **[0002]** In some other aspects, the present technology relates to methods associated with monetizing human attention (and attention from any intelligent entity) more effectively than existing online advertising model which is suffering from diminishing returns on incremental advertisement (ad, ads) targeting improvements.

15 **[0003]** In yet other aspects, activities that are described in this patent disclosure as happening on an external network in which multiple intelligent entities participate in collaborative problem-solving, can also be implemented within a single computerized intelligent system where the intelligent entities are all computerized or AI agents that reside within that single computerized intelligent system.

2.0 BACKGROUND ART

20 **[0004]** The fastest and safest path to development of AGI and SI has been described in previous invention disclosures. Methods and catalysts for increasing intelligence of AI systems generally, as well as the development of AGI and Personalized SuperIntelligence (PSI) have also been previously disclosed. Therefore, the following U.S. Provisional Patent Applications (PPA), are incorporated herein by reference.

25 **[0005]** The present application incorporates by reference all work in the PPA No. 63/487,494 entitled: Advanced Autonomous Artificial Intelligence (AAAI) System and Methods, which was filed and received by the USPTO on February 28, 2023.

30 **[0006]** The present application incorporates by reference all work in the PPA No. 63/491,040 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.

[0007] The present application incorporates by reference all work in the PPA No. 63/577,830 entitled: System and Methods for Human-Centered AGI, which was filed with the USPTO on March 24, 2023.

5 [0008] The present application incorporates by reference all work in the PPA No. 63/628,410 entitled: System and Methods for Safe, Scalable, Artificial General Intelligence, which was filed with the USPTO on July 18, 2023.

[0009] The present application incorporates by reference all work in the PPA No. 63/519,549 entitled: Safe Personalized Super Intelligence (PSI), which was filed with the USPTO on August 14, 2023.

10 [0010] The present application incorporates by reference all work in the PPA No. 63/601,930 entitled: Catalysts for Growth of SuperIntelligence, which was filed with the USPTO on November 22, 2023.

[0011] The present application incorporates by reference all work in the PPA No. 63/609,800 entitled: System and Methods for Safe Alignment of SuperIntelligence, which was filed with the
15 USPTO on December 13, 2023.

[0012] In addition to the above-mentioned PPAs, the present application 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
20 (filed on 2/26/2024); PCT/US24/19486 (filed on 3/12/2024); and PCT/US24/20334 (filed on 3/17/2024).

[0013] The present application contains further technologies that can be used with the system and methods described in the above-mentioned PPAs and PCTs as well as in a standalone fashion.

[0014] While the above-described devices fulfill their respective, particular objectives and
25 requirements, the aforementioned patents do not describe an online advertising technology for AGI and SI that allows increasing the intelligence of AI/AGI systems by overcoming the data bottleneck that AI researchers currently face and/or monetizes human attention more effectively than existing online advertising model which is suffering from diminishing returns on incremental ad, ads targeting improvements.

30 [0015] Therefore, a need exists for a new and improved online advertising technology for AGI and SI that can be used for increasing the intelligence of AI/AGI systems by overcoming the data bottleneck that AI researchers currently face and that monetizes human attention more effectively than existing online advertising model, and that enables ad-targeting improvements. In this regard, the present technology substantially fulfills this need. In this respect, the online advertising

technology for AGI and SI according to the present technology substantially departs from the conventional concepts and designs of the prior art, and in doing so provides an apparatus primarily developed for the purpose of increasing the intelligence of AI/AGI systems by overcoming the data bottleneck that AI researchers currently face. The technology also monetizes human attention more effectively than existing online advertising model and enables ad-targeting improvements.

DISCLOSURE OF TECHNOLOGY

[0016] In view of the foregoing disadvantages inherent in the known types of AI systems and methods, and in existing computerized advertisement monetization methods, at least some embodiments of the present technology provide a novel online advertising technology for AGI and SI, and overcome one or more of the mentioned disadvantages and drawbacks of the prior art. As such, the general purpose of at least some embodiments of the present technology, which will be described subsequently in greater detail, is to provide a new and novel online advertising technology that powers AGI and SI and which has all the advantages of the prior art mentioned herein as well as many novel features. The technology results in an online advertising technology for AGI and SI which is not anticipated, rendered obvious, suggested, or even implied by the prior art, either alone or in any combination thereof.

[0017] According to one aspect, the present technology can include a system or method of utilizing an interactive online advertisement to train or increase knowledge of an Artificial Intelligence (AI) agent or system. An interactive online advertisement is populated utilizing problem information relating to a problem provided to or by an AI agent or system. The interactive online advertisement is configured or configurable to acquire human attention from human (or other intelligent entity) users with content associated with the problem information or the problem. Input information acquired by the interactive online advertisement from the human (or intelligent entity) users is received by the AI agent or system. The AI agent or system utilizes the input information to update the interactive online advertisement or provides the input information for use in a problem solving process on the problem.

[0018] According to another aspect, the present technology can include a System for utilizing online advertising technology for increasing an intelligence of an AI agent or system. The inventive System can include:

a computer system comprising: a processor, a computer-readable storage medium, and program instructions stored on the computer-readable storage medium being executable by the processor to cause the computer system to:

provide the online advertisement unit to one or more user computer systems utilizing a network;

5 receive one or more inputs from the user computer systems participating in the online advertisement unit, the inputs being providing by a human user utilizing the user computer systems respective, and the inputs being associated with any one of or any combination of solving the problem, solving a sub-problem of the problem, and advancing progress on the problem;

communicate the inputs to the AI agent or system;

10 utilize the inputs by the AI agent or system in a universal problem solving method to generate a solution to the problem or the sub-problem; and

train the AI agent or system with results from the universal problem solving method that are based on successful or unsuccessful solution attempts to solve the problem or the sub-problem, thereby increasing an intelligence of the AI agent or system.

[0019] According to yet another aspect, the present technology can include a method for online advertising technology utilized with an AI agent or system for increasing an intelligence of the AI agent or system. The method can include the steps of:

receiving problem information of a problem provided by an AI agent or system;

populating an online advertisement unit including the problem information;

20 providing the online advertisement unit to one or more user computer systems utilizing a network;

receiving one or more inputs from the user computer systems participating in the online advertisement unit, the inputs being providing by a human user utilizing the user computer systems respective, and the inputs being associated with any one of or any combination of solving the problem, solving a sub-problem of the problem, and advancing progress on the problem;

25 communicating the inputs to the AI agent or system;

utilizing the inputs by the AI agent or system in a universal problem solving method to generate a solution to the problem or the sub-problem; and

30 training the AI agent or system with results from the universal problem solving method that are based on successful or unsuccessful solution attempts to solve the problem or the sub-problem, thereby increasing an intelligence of the AI agent or system.

[0020] Some embodiments of the present technology can include a step of receiving advertisement specifications from a (human or intelligent entity) client to be used in the populating the online advertisement unit in combination with the problem information.

[0021] In some embodiments, the advertisement specifications can include any one of or any combination of advertisement content, demographic information, location restrictions for the online advertisement unit, advertisement budget, and metrics for determining a successful solving of the problem or the sub-problem.

- 5 [0022] Some embodiments of the present technology can include a step of billing the client based on cost per thousand impressions or clickthrough rate metrics.

[0023] In some embodiments, the problem information can include any one of or any combination of a current problem state of the problem, a sub-problem of the problem, a goal of the problem, and a sub-goal of the goal.

- 10 [0024] Some embodiments of the present technology can include a step of communicating to the AI agent or system one or more additional inputs from intelligent entities in combination with the inputs from the participating user computer systems, the intelligent entities being any one of or any combination of an additional human user utilizing an additional computer system, an additional AI agent or system, an Artificial General Intelligent (AGI) agent or system, and a SuperIntelligent (SI)
15 agent or system.

- [0025] In some embodiments, the updating of the online advertisement unit can include any one of or any combination of whether the inputs was accepted or rejected, whether a goal or sub-goal of the problem has been achieved, operators or list(s) of potential operators, if the input was selected by the AI agent or system to advance problem solving, a description of a new problem state after
20 applying the selected inputs, an amount of credit or payment that the intelligent entities has accumulated based on inputs from the intelligent entities, and an additional request for new user input by way of the online Ad unit or a link that directs the intelligent entities to an interface for participating in the problem solving process on the problem.

- [0026] Some embodiments of the present technology can include a step of crediting the intelligent
25 entities if the problem or the sub-problem is solved, based on an amount of contribution by the participating intelligent entities.

[0027] Some embodiments of the present technology can include a step of soliciting, if the problem or the sub-problem is not solved, additional input from the human user of the participating user computer systems respectively.

- 30 [0028] Some embodiments of the present technology can include a step of compensating the human user when the human user exits the online Ad advertisement unit and stops participating in the problem solving process.

[0029] In some embodiments, the training of the AI agent or system can further include any one of or any combination of training methods selected from the group consisting of differential privacy,

federated learning, homomorphic encryption, synthetic data generation, secure multi-party computation, data anonymization, transfer learning, active learning, self-supervised learning, domain adaptation, reinforcement learning, few-shot learning, explainable AI, privacy-preserving record linkage, data augmentation, generative adversarial networks, crowdsourcing for data labelling, model personalization layers, knowledge distillation, and ensemble learning.

[0030] Some embodiments of the present technology can include a step of creating a database of intelligent entity experts, wherein the intelligent entity experts are any one of or any combination of human user utilizing a computer system, an additional AI agent or system, an Artificial General Intelligent (AGI) agent or system, and a SuperIntelligent (SI) agent or system.

[0031] In some embodiments, the online advertisement unit can be provided to one or more of the intelligent entity experts from the database.

[0032] In some embodiments, the creating of the database can utilize a technology selected from the group consisting of relational database management system, NoSQL databases, data warehousing, data mining, machine learning algorithms, graph databases, vector databases, indexing, full-text search, blockchain, data visualization tools, application programming interfaces, data cleaning and preprocessing, cloud storage, caching, transactional database systems, real-time database systems, content delivery network, data compression, data encryption, and replication.

[0033] Some embodiments of the present technology can include a step of bidding on attention using an attention spot market.

[0034] In some embodiments, the attention spot market can include any one of or any combination of:

- a means for the human (or intelligent entity) users with attention to sell to access a marketplace and specify seller information for sale and an ask price for the information;
- a means for buyers of human (or intelligent entity) attention to access the marketplace and specify buyer information the buyer is willing to buy and a bid price for buying the buyer information;
- a market mechanism for queuing the bid price and the ask price, including categories of the buyer and seller information, wherein each category has a market in the marketplace;
- the market mechanism is configured or configurable to make the market in each category by market makers; and
- the market mechanism is configured or configurable to match bid prices and ask prices, and a transaction occurs that is binding on the buyer and the seller of the information.

[0035] In some embodiments, the attention spot market can include the steps of:

buyers and sellers of human (or intelligent entity) attention and expertise register on a platform, the buyers providing details about buyer interests or expertise, the sellers providing details about sellers interests or expertise;
listing by the sellers available time slots and expertise areas, and listing by the buyers needs and
5 time slots the buyers is interested in;
utilizing, by the platform, an algorithm to dynamically price human (or intelligent entity) attention and expertise based on supply, demand, and user ratings;
matching one or more of the buyers and one or more of the sellers based on requirements, availability, and price;
10 enabling transactions where the buyers pay for the seller time slots, and wherein the platform takes a commission; and
providing feedback where after each session, the buyers and the sellers rate each other, influencing future pricing and matching.

[0036] In some embodiments, the platform can utilize any one or any combination of mechanisms
15 selected from the group consisting of Dutch auction, reverse auction, sealed bid auction, open ascending price auction, fixed price with time priority, dynamic pricing based on ratings, supply-demand curve adjustment, time slot segmentation, subscription access, freemium model, group buying, tiered expertise levels, flash sales, loyalty points system, hybrid auction, geographic pricing, behavioral pricing, escrow system, social influence pricing, and tokenized transactions.

20 [0037] In some embodiments, the attention spot market can include the steps of:
buyers and sellers of human (or intelligent entity) attention and expertise register on a platform, the buyers providing details about buyer interests or expertise, the sellers providing details about sellers interests or expertise;
creating an auction by the sellers for seller time slots;
25 placing one or more buyer bids by a buyer on time slots and expertise the buyer requires;
closing the auction at a predetermined time or when the seller accepts a buyer bid;
paying the seller by the buyer of a winning bid, and the seller provides the attention or expertise based on the time slot, wherein the platform mediates an exchange and secures payment; and
providing feedback where after each auction, the buyers and the sellers rate each other, affecting
30 future auctions and visibility on the platform.

[0038] In some embodiments, the platform can utilize any one of or any combination of mechanisms selected from the group consisting of Dutch auction, reverse auction, sealed bid auction, open ascending price auction, fixed price with time priority, dynamic pricing based on ratings, supply-demand curve adjustment, time slot segmentation, subscription access, freemium

model, group buying, tiered expertise levels, flash sales, loyalty points system, hybrid auction, geographic pricing, behavioral pricing, escrow system, social influence pricing, and tokenized transactions.

5 [0039] In some embodiments, the online advertisement unit can be configured or configurable to include material relevant to an area of expertise associated with the problem information.

[0040] In some embodiments, the online advertisement unit can be configured or configurable to, upon interaction by the human (or intelligent entity) users, expand or redirect the human (or intelligent entity) user to a secure form on a landing page, wherein the form is configured to collect user information and to provide questions to the human (or intelligent entity) user regarding to
10 qualifications of the human (or intelligent entity) user.

[0041] Some embodiments of the present technology can include a step of validating the user information and categorizing the human (or intelligent entity) users in a database based on the qualifications.

[0042] In some embodiments, the online advertisement unit can be configured or configurable to
15 include an interface allowing the human (or intelligent entity) users to input a solution to the problem or the sub-problem.

[0043] In some embodiments, the online advertisement unit can be configured or configurable to include any one of or any combination of elements of gamification, a scoring system, and a leaderboard of other human (or intelligent entity) users.

20 [0044] In some embodiments, the online advertisement unit can be configured or configurable to include text input fields or other interactive tools enabling the human (or intelligent entity) users to provide a solution to the problem or sub-problem directly within a space of the online advertisement unit.

[0045] In some embodiments, the solution can be stored in a database, and an algorithm is utilized
25 to assess a quality and relevance of the solution based on predefined criteria.

[0046] In some embodiments, the human (or intelligent entity) users can receive instant feedback or rewards points based on the assessed quality and relevance of the solution, redeemable for various incentives.

[0047] In some embodiments, contributions by the human (or intelligent entity) users can be tracked
30 to identify top contributors from multiple additional human (or intelligent entity) users for potential future engagement.

[0048] In some embodiments, the online advertisement unit can be configured or configurable to, upon interaction by the human (or intelligent entity) users, would expand or redirect the human (or intelligent entity) user to a secure form on a landing page, wherein the form is configured to collect

user information and to provide questions to the human (or intelligent entity) user regarding to qualifications of the human (or intelligent entity) user.

5 [0049] Some embodiments of the present technology can include a step of validating the user information and categorizing the human (or intelligent entity) users in the database based on the qualifications.

[0050] Some embodiments of the present technology can include steps of:
analyzing the inputs from the participating human (or intelligent entity) users to determine a
quality or relevance of the inputs, respectively; and
compensating the participating human (or intelligent entity) users based on the quality or
10 relevance of the inputs.

[0051] Some embodiments of the present technology can include a step of providing digital wallet information by each of the participating human (or intelligent entity) users into a secure form, wherein the secure form is within the online advertisement unit or is provided by a linked platform.

15 [0052] In some embodiments, the compensation can be provided in any one of or any combination of cryptocurrency, electronic gift cards, electronic vouchers, subscriptions, access to online resources or content, recognition on online platforms, merchandise related to an expertise of the human (or intelligent entity) users respectively, and sponsorships for professional conferences or events.

20 [0053] Some embodiments of the present technology can include a step of incentivizing the participating human (or intelligent entity) users to refer other human (or intelligent entity) users to an online platform, by offering additional compensation for each referral that contributes to the platform, wherein the online platform receives the problem and provides the online advertisement unit.

25 [0054] Some embodiments of the present technology can include a step of receiving reputational metrics on each of the human (or intelligent entity) users.

[0055] In some embodiments, the reputational metrics can be multi-dimensional including any one of or any combination of timeliness, budget compliance, quality of work, solution success rate, client satisfaction, peer-rating of competency, external reputation, innovation score, communication skills, adaptability, leadership quality, technical proficiency, learning agility, conflict resolution, project
30 management skills, reliability, efficiency, cultural fit, work ethic, client retention rate, feedback responsiveness, networking ability, mentorship and training, availability and responsiveness, problem-solving speed, and creativity and innovation.

[0056] In some embodiments, the reputational metrics for any one of the human (or intelligent entity) users can be automatically updated within the online advertisement unit.

[0057] Some embodiments of the present technology can include a step of updating the online advertisement unit based on the results from the universal problem solving method conducted on the problem or the sub-problem.

5 [0058] In some embodiments, the online advertisement unit can be configured or configurable to ask the human (or intelligent entity) users to provide a description of how the problem is to be framed, and then the description is then translated into a language of the universal problem solving method using a natural language to problem solving language translator.

[0059] In some embodiments, the online advertisement unit can be configured or configurable to ask the human (or intelligent entity) users to provide a description of intermediate goals towards a final
10 solution of the problem, and wherein the intermediate goals are communicated to the AI agent or system for processing by the universal problem solving method.

[0060] Some embodiments of the present technology can include a step of requesting, by the AI agent or system, the human (or intelligent entity) users to provide safety or ethical information associated with the problem, the sub-problem or goals associated with the problem.

15 [0061] Some embodiments of the present technology can include a step of utilizing the safety or ethical information in the populating or in updating the online advertisement unit.

[0062] Some embodiments of the present technology can include a step of modifying the online advertisement unit to ask the human (or intelligent entity) users to vote, rank or rate one or more potential goals, sub-goals or actions that have been proposed by any one of or any combination of
20 the human users, and intelligent entities.

[0063] Some embodiments of the present technology can include a step of requiring the AI agent or system to be identified as an AI agent so any inputs from the AI agent or system is excluded where only human opinions are sought.

25 [0064] In some embodiments, the online advertisement can be is provided by an online advertisement service provider that uses user information about the human (or intelligent entity) users to target the human(or intelligent entity) users with expertise relevant to the problem.

[0065] Some embodiments of the present technology can include a step of measuring and recording response rates for the participating human (or intelligent entity) users who contribute information within the online advertisement unit or by way of links to websites or interfaces outside the online
30 advertisement unit.

[0066] Some embodiments of the present technology can include a step of recording metrics related to a quality of information or solutions provided by the targeted human (or intelligent entity) users.

[0067] Some embodiments of the present technology can include a step of calculating correlations or statistical relationships on the user information to determine any one of or any combination of

which factors have a highest impact on the metrics, which of the metrics is prioritized by the online advertisement service provider, and which of the inputs was most effective at meeting or maximizing problem solving or training of the AI agent or system.

5 [0068] Some embodiments of the present technology can include a step of performing arbitrage that can include the steps of:

estimating, by the online advertisement service provider, a value of a particular type of human (or intelligent entity) attention and knowledge with regard to the problem or sub-problem; estimating, by the online advertisement service provider, a cost of obtaining an amount of the human(or intelligent entity) attention required to complete the problem or sub-problem from
10 the human users or intelligent entities by way of the online advertisement unit; and determining if the estimated value exceeds the estimated cost by a predetermined or dynamic variable representing a profit margin, then purchase the human (or intelligent entity) attention at cost approximating that of the estimated value and sell the human (or intelligent entity) attention to clients or use the human (or intelligent entity) attention to perform tasks
15 that create value approximating that of the estimated cost.

[0069] Some embodiments of the present technology can include a step of performing arbitrage in priority order of a largest arbitrage opportunities first, until a minimal acceptable arbitrage opportunity is reached.

20 [0070] In some embodiments, the online advertisement unit can include one or more interfaces selected from the group consisting of any one of or combination of Large Language Model interface, Small Language Model interface, natural language text interface, audio-based interfaces, text boxes, dropdown lists, templates, dynamically sizing input and output areas, visual input or output devices, virtual reality devices, multimodal interfaces, and augmented reality devices.

25 [0071] In some embodiments, the online advertisement unit can be multiple online advertisement units with a first online advertisement unit being provided to a first set of the human (or intelligent entity) users for receipt of the inputs, and a second online advertisement unit being provided to the first set or a second set of the human (or intelligent entity) users and asking to vote on the inputs.

[0072] In some embodiments, the inputs can be provided to the AI agent or system in real-time for real-time utilization by the universal problem solving method.

30 [0073] According to still another aspect, the present technology can include a method of increasing monetization of online advertising revenue for an online advertisement service provider by delivering to a client information based on human (or intelligent entity) attention. The method can include the steps of:

building, by an online advertisement service provider, one or more databases containing information about human (or intelligent entity) users, the information is related to targeting specific types of online advertisements to specific human (or intelligent entity) users; communicating to the online advertisement service provider by a client a request to purchase human (or intelligent entity) attention, the request including requirements and desired categories of the human (or intelligent entity) attention; performing checks on the request by the online advertisement service provider at a time that the requirements are communicated, the checks are configured or configurable to ensure that the request does not violate regulations or ethical requirements; creating content for an interactive online advertisement unit that is configured or configurable to capture human (or intelligent entity) attention and work from human (or intelligent entity) users; deploying, by the online advertisement service provider, the interactive online advertisement unit to a computer system of the human (or intelligent entity) users by way of a network; allowing the interactive online advertisement unit to collaborate with multiple of the human (or intelligent entity) users and integrate with real-time or asynchronous data capabilities; performing ethics and safety checks during a problem solving process on a problem provided by the client, the ethics and safety checks are configured or configurable to ensure that unethical or unsafe expertise is not part of the problem; and improving the interactive online advertisement unit based on a feedback loop including on any one of or any combination of reputational metrics, and metrics related to the problem solving process or knowledge captured.

[0074] In some embodiments, the client can be any one of or any combination of an AI agent or system, a system that trains or customizes AI agents, and a system that provides solutions to problems.

[0075] Some embodiments of the present technology can include a step of determining if the database does not include a predetermined number of the human (or intelligent entity) users, or if additional human (or intelligent entity) users are required, then configuring the interactive online advertisement unit to acquire additional human (or intelligent entity) users.

[0076] In some embodiments, the requirements can be communicated by way of interaction with a human (or intelligent entity) attention spot market.

[0077] In some embodiments, the attention spot market can include any one of or any combination of:

a means for the human (or intelligent entity) users with attention to sell to access a marketplace and specify seller information for sale and an ask price for the information;

a means for buyers of human (or intelligent entity) attention to access the marketplace and specify buyer information the buyer is willing to buy and a bid price for buying the buyer information;

a market mechanism for queuing the bid price and the ask price, including categories of the buyer and seller information, wherein each category has a market in the marketplace;

the market mechanism is configured or configurable to make the market in each category by market makers; and

the market mechanism is configured or configurable to match bid prices and ask prices, and a transaction occurs that is binding on the buyer and the seller of the information.

[0078] In some embodiments, the attention spot market can include the steps of:

buyers and sellers of human (or intelligent entity) attention and expertise register on a platform, the buyers providing details about buyer interests or expertise, the sellers providing details about sellers interests or expertise;

listing by the sellers available time slots and expertise areas, and listing by the buyers needs and time slots the buyers is interested in;

utilizing, by the platform, an algorithm to dynamically price human (or intelligent entity) attention and expertise based on supply, demand, and user ratings;

matching one or more of the buyers and one or more of the sellers based on requirements, availability, and price;

enabling transactions where the buyers pay for the seller time slots, and wherein the platform takes a commission; and

providing feedback where after each session, the buyers and the sellers rate each other, influencing future pricing and matching.

[0079] In some embodiments, the platform can utilize any one or any combination of mechanisms selected from the group consisting of Dutch auction, reverse auction, sealed bid auction, open ascending price auction, fixed price with time priority, dynamic pricing based on ratings, supply-demand curve adjustment, time slot segmentation, subscription access, freemium model, group buying, tiered expertise levels, flash sales, loyalty points system, hybrid auction, geographic pricing, behavioral pricing, escrow system, social influence pricing, and tokenized transactions.

[0080] In some embodiments, the attention spot market can include the steps of:

buyers and sellers of human (or intelligent entity) attention and expertise register on a platform,
the buyers providing details about buyer interests or expertise, the sellers providing details
about sellers interests or expertise;
creating an auction by the sellers for seller time slots;
5 placing one or more buyer bids by a buyer on time slots and expertise the buyer requires;
closing the auction at a predetermined time or when the seller accepts a buyer bid;
paying the seller by the buyer of a winning bid, and the seller provides the attention or expertise
based on the time slot, wherein the platform mediates an exchange and secures payment; and
providing feedback where after each auction, the buyers and the sellers rate each other, affecting
10 future auctions and visibility on the platform.

[0081] In some embodiments, the platform can utilize any one of or any combination of
mechanisms selected from the group consisting of Dutch auction, reverse auction, sealed bid
auction, open ascending price auction, fixed price with time priority, dynamic pricing based on
ratings, supply-demand curve adjustment, time slot segmentation, subscription access, freemium
15 model, group buying, tiered expertise levels, flash sales, loyalty points system, hybrid auction,
geographic pricing, behavioral pricing, escrow system, social influence pricing, and tokenized
transactions.

[0082] Some embodiments of the present technology can include a step of providing a feedback
mechanism associated the human (or intelligent entity) attention spot market, the feedback
20 mechanism is configured or configurable to record each and every step in the problem solving
process, and to create a vector track record of a performance of all the human (or intelligent entity)
users on every task and sub-task.

[0083] In some embodiments, the vector track record can be implemented by way of blockchain
technology to allow for precise reputations that are analyzed by an AI agent or system and converted
25 into estimates of a value for each of the human (or intelligent entity) users for any of the tasks.

[0084] In some embodiments, the interactive online advertisement unit can be configured or
configurable to capture the human (or intelligent entity) attention and work from within the
interactive online advertisement unit or by providing link to one or more webpages or interfaces that
are optimized for the capture outside of the interactive online advertisement unit.

30 [0085] In some embodiments, the building of the databases can include the step of purchasing, by
the online advertisement service provider, user data and information from remote sources.

[0086] According to yet still another aspect, the present technology can include a method for online
advertising technology utilized with an AI agent or system. The method can include the steps of:

populating, by an AI agent or system, an online advertisement unit including problem information provided to the AI agent or system;

providing the online advertisement unit to one or more user computer systems utilizing a network;

5 receiving one or more inputs from the user computer systems participating in the online advertisement unit, the inputs being providing by a human (or intelligent entity) user utilizing the user computer systems respective, and the inputs being associated with any one of or any combination of solving the problem, solving a sub-problem of the problem, and advancing progress on the problem;

10 communicating the inputs to the AI agent or system;

utilizing the inputs by the AI agent or system or an additional AI agent or system in a universal problem solving method to generate a solution to the problem or the sub-problem; and

utilizing a feedback mechanism that provides metrics to the AI agent or system to train the AI agent or system, the metrics assign credit or blame to the participating human (or intelligent entity) users and targeting mechanisms that directs the online advertisement unit to
15 particular human (or intelligent entity) users.

[0087] Some embodiments of the present technology can include a step of measuring and recording response rates for the participating human (or intelligent entity) users who contribute information within the online advertisement unit or by way of links to websites or interfaces outside the online
20 advertisement unit.

[0088] Some embodiments of the present technology can include a step of recording the metrics related to a quality of information or solutions provided by the targeted human (or intelligent entity) users.

[0089] Some embodiments of the present technology can include a step of calculating correlations or statistical relationships on the user information to determine any one of or any combination of
25 which factors have a highest impact on the metrics, which of the metrics is prioritized by the online advertisement service provider, and which of the inputs was most effective at meeting or maximizing problem solving or training of the AI agent or system.

[0090] Some embodiments of the present technology can include a step of performing arbitrage that
30 can include the steps of:

estimating, by the online advertisement service provider, a value of a particular type of human (or intelligent entity) attention and knowledge with regard to the problem or sub-problem;

estimating, by the online advertisement service provider, a cost of obtaining an amount of the human (or intelligent entity) attention required to complete the problem or sub-problem from the human users or intelligent entities by way of the online advertisement unit; and determining if the estimated value exceeds the estimated cost by a predetermined or dynamic variable representing a profit margin, then purchase the human (or intelligent entity) attention at cost approximating that of the estimated value and sell the human (or intelligent entity) attention to clients or use the human (or intelligent entity) attention to perform tasks that create value approximating that of the estimated cost.

[0091] Some embodiments of the present technology can include a step of performing arbitrage in priority order of a largest arbitrage opportunities first, until a minimal acceptable arbitrage opportunity is reached.

[0092] There has thus been outlined, rather broadly, features of the present technology in order that the detailed description thereof that follows may be better understood and in order that the present contribution to the art may be better appreciated.

[0093] Numerous objects, features and advantages of the present technology will be readily apparent to those of ordinary skill in the art upon a reading of the following detailed description of the present technology, and illustrative embodiments of the present technology, taken in conjunction with the accompanying drawings.

[0094] As such, those skilled in the art will appreciate that the conception, upon which this disclosure is based, may readily be utilized as a basis for the designing of other structures, methods and systems for carrying out the several purposes of the present technology.

[0095] It is therefore an object of the present technology to provide a new and novel online advertising technology for AGI and SI that has all of the advantages of the prior art AI systems and methods, and computerized advertisement monetization methods and none of the disadvantages.

[0096] It is another object of the present technology to provide a new and novel online advertising technology for AGI and SI that may be easily and efficiently implemented and marketed.

[0097] An even further object of the present technology is to provide a new and novel online advertising technology for AGI and SI that has a low cost of implementation with regard to both resources and labor, and which accordingly is then susceptible of low prices of sale to the consuming public, thereby making such online advertising technology for AGI and SI economically available to the buying public.

[0098] Still another object of the present technology is to provide a new online advertising technology for AGI and SI that provides in the system and methods of the prior art some of the

advantages thereof, while simultaneously overcoming some of the disadvantages normally associated therewith.

[0099] For a better understanding of the present technology, its operating advantages and the specific objects attained by its uses, reference should be made to the accompanying drawings and descriptive matter in which there are illustrated embodiments of the present technology. Whilst multiple objects of the present technology have been identified herein, it will be understood that the following description is not limited to meeting most or all of the objects identified and that some embodiments of the present technology may meet only one such object or none at all.

10 BRIEF DESCRIPTION OF THE DRAWINGS

[00100] The technology will be better understood and objects other than those set forth above will become apparent when consideration is given to the following detailed description thereof. Such description makes reference to the annexed drawings wherein:

15 [00101] FIG. 1 is a flow chart illustrating an embodiment of the subsystems utilizable in the AAAI system and method of the present technology.

[00102] FIG. 2 is a block diagram illustrating an exemplary process of the overall process utilizable with the present technology.

20 [00103] FIG. 3 is a flow chart illustrating an exemplary embodiment of the system and methods for creating an ethical and safe AGI or PSI from the collective intelligence of AAAIs and humans utilizable with the present technology.

[00104] FIG. 4 is a flow chart illustrating an exemplary embodiment of the scalable universal problem solving architecture including system and methods for human-centered AGI, with relevance for PSIs, constructed in accordance with the principles of the present technology.

25 [00105] FIG. 5 is a flow chart illustrating an exemplary embodiment of the scalable solution learning subsystem or process.

[00106] FIG. 6 is a flow chart illustrating an exemplary embodiment of the scalable natural language to problem solving language translator subsystem or process.

[00107] 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.

30 [00108] FIG. 8 is a flow chart illustrating an exemplary embodiment of the scalable safety and ethics checks subsystem or process, wherein AIs or AAAIs might also be PSIs.

[00109] FIG. 9 is a diagram illustrating various use cases for domain-specific problems which depend upon the underlying WorldThink protocol, and which together help form the basis for an AGI system capable of solving a wide range of problems, wherein the AAAIs identified in the

Figure might also be PSIs.

[00110] FIG. 10 is a diagram illustrating some of the steps in the universal problem solving framework that is part of the WorldThink protocol and used by the AAAI system and which also can be used by PSIs or intelligent entities.

- 5 [00111] FIG. 11 is a flow chart illustrating some of the basic problem solving functionality supported by the invention utilizable with the system and method of the present technology, and wherein problem solvers might be humans, AIs, PSIs, or intelligent entities.

[00112] FIG. 12 is a flow chart illustrating some of the basic problem solving functionality supported by the current invention utilizing multiple problem solvers, which could be humans, AIs,

- 10 PSIs, or intelligent entities collaborating to solve a client problem.

[00113] FIG. 13 is a diagram illustrating features and functions of the Problem Solving Tree structure in the problem solving / WorldThink protocol of the AGI system that is utilizable with the present technology.

[00114] FIG. 14 is a diagram illustrating the current technology for Online Advertising.

- 15 [00115] FIG. 15 is a diagram illustrating the basic components of the Attention Spot Market, which is part of the present inventive technology.

[00116] FIG. 16 is a diagram illustrating the Direct Exchange Platform Implementation variation of the Attention Spot Market technology.

- 20 [00117] FIG. 17 is a diagram illustrating the Auction-Based Marketplace Implementation variation of the Attention Spot Market technology.

[00118] FIG. 18 is a diagram illustrating an exemplary inventive method for implementing problem solving within an Ad unit.

[00119] FIG. 19 is a diagram illustrating an inventive basic feedback process for Ad targeting.

- 25 [00120] FIG. 20 is a diagram illustrating an exemplary feedback process for the attention spot market component of the present technology.

[00121] FIG. 21 is a schematic block diagram illustrating an exemplary electronic computing device that may be used to implement an embodiment of the present technology.

[00122] The same reference numerals refer to the same parts throughout the various figures.

30 DETAILED DESCRIPTION OF THE TECHNOLOGY

3.0 Definitions

[00123] Artificial Intelligence (AI) – A non-human entity capable of behavior that most humans would consider intelligent in at least one area, or in some respect.

[00124] Artificial General Intelligence (AGI) – Conventionally refers to an AI that is capable of doing all (or almost all) intellectual tasks that 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 SuperIntelligent AGI that can do all intellectual task as well or 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 described as being implemented by a system and associated methods.

[00125] 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 useful 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.

[00126] 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.

[00127] AI Ethics – The ethics adopted by an AI or AGI that describe what is right and wrong in given contexts.

[00128] Alignment Problem – The problem that arises when AI Ethics are not aligned with Human Ethics resulting in AI or AGI taking actions that humans consider unethical and/or which are dangerous to individual humans or the human race.

[00129] 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 with information for specific tasks.

[00130] Collective Intelligence (CI) – The intelligence that emerges when multiple intelligent entities are focused on solving a common problem, or when the knowledge from multiple intelligent entities is pooled to overcome limits of bounded rationality. Collective Intelligence historically has been human collective intelligence, but AGI is based on collective intelligence of both 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 steps that are useful 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, the humans are exhibiting 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, and then uses that analysis to direct targeted ads to the human(s).

5 [00131] 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
10 is wrong must be accepted as premises, from which starting point an intelligent entity can propose rational actions to realize those values or ethics.

[00132] 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,
15 misleading or false.

[00133] Human Ethics – The ethics asserted by human beings which describe what is right and wrong in given contexts.

[00134] Intelligent Entities or Entity – A human utilizing a computer system, an AI agent or system including AGI and SI systems, a clone of an AI agent or system, an AAAI agent or system, and/or a
20 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 a problem, 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.

25 [00135] Large Language Model (LLM) – A type of AI that can accept natural language as an input and generate natural language as an output. Typically, LLMs are trained using ML techniques on large datasets so that they can 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
30 representations and input and generate language and/or image and/or visual representations as output. For the purposes of 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 the input or the output. LLMs can also act as a type of AI agent and are sometimes referred to as such in the present technology. For purpose of this disclosure, Small Language Models (SLMs) are also included in the definition of LLM.

- [00136] Machine Learning (ML) – A sub-field that is concerned with developing AI by enabling machines to teach themselves or learn their knowledge rather than such knowledge being explicitly programmed into them (as would be the case with an Expert System AI developed via classical knowledge engineering methods).
- 5 [00137] Narrow AI – An AI that performs at human or at super-human levels in a relatively restricted domain such as game playing, brewing beer, analyzing legal contracts, etc. Narrow AI is contrasted with AGI that can perform at human level at ALL intellectual tasks. 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.
- 10 [00138] 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.
- [00139] Prohibited Attributes – Requests, goals, problems, terms, phrases, questions, answers, solutions, information and the like that are determined or set as being illegal, immoral, unethical,
- 15 dangerous, deadly and the like. For example, requesting information for getting Molotov Cocktails through airport security.
- [00140] Safety – Generally, the concern for human safety and survival is distinct from ethics and values.
- [00141] Safety Feature – An aspect of the design or operation of the present technology which
- 20 increases the safety of one or more humans, often by helping increase the probability that AI ethics align with human ethics, thus surmounting the Alignment Problem.
- [00142] Training/Tuning/Customization – Conventionally the term “training” is used to denote training a network (e.g., LLM) to behave intelligently. Tuning refers to activities that fine-tune the trained base model so that it performs even better, typically at specific tasks. Customizing refers to a
- 25 wide variety of activities including, but not limited to, training and tuning that make an AI uniquely suited for the purposes of 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, and the degree and type of effort involved varies, the aim of all three is to adapt the AI and make it behave more intelligently or more uniquely suited to a particular user(s) or
- 30 application(s).
- [00143] Weights/Weights of the Network – In the field of machine learning, many systems learn by adjusting the weights in a neural network architecture that can be represented as a network of nodes and links between nodes. The weight of a link connecting two nodes, for example, may correspond to the strength of association or connection between the whatever nodes represent. These weights

can also represent excitatory or inhibitory connections between concepts, as in a neural network representation. 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. By combining, manipulating, or otherwise changing this numerical information, the learning, knowledge, or expertise and behavior of the system can be changed.

10

4.0 Background for the Present Technology

[00144] In the following description, for purposes of explanation and not limitation, specific details are set forth, such as particular embodiments, procedures, techniques, etc. in order to provide a thorough understanding of the present technology. However, it will be apparent to one skilled in the art that the present technology may be practiced in other embodiments that depart from these specific details.

15

[00145] It can be appreciated that the present technology provides a technical effect, contribution and solution with a technical implementation of multiple customized AAAI systems communicating over a collective intelligence network, in combination with all the AAAI systems each utilizing a common cognitive architecture including one or more problem solving protocols for generating one or more solutions or answers to a problem request, and providing the solutions or answers to a user for approval. Where the customization of the AI system resulting in the AAAI includes input from human users for training the AI or the AAAI. Further technical contribution or solution can be where the multiple customized AAAI systems can include one or more cloned AAAIs that can each be customized independently of a parent AAAI and independent of other cloned AAAIs of the same system.

20

[00146] Still another technical contribution and solution is for the faster and safer creating of AGI that utilizes human input in training and customization for imparting human ethical attributes to the AAAI and/or AGI.

[00147] Yet still another technical contribution and solution is for providing an online advertising technology for AI, AGI and SI for use in connection with increasing the intelligence of AI/AGI/SI agents or systems by overcoming the data bottleneck that AI researchers currently face. Still further, a technical contribution and solution is for monetizing human attention (and attention from any

30

intelligent entity) more effectively than existing online advertising model which is suffering from diminishing returns on incremental advertisement (ad, ads) targeting improvements.

[00148] Still yet another technical contribution and solution is for providing improved solutions or answers to a user's problem request that have a higher chance of acceptance by the user as the provided solutions or answers will have been generated by AAAs with similar training to the user's AAAI thereby aligning with the user's parameters.

[00149] It can be appreciated that the present technology is found outside of computer program exclusion and/or abstract idea interpretation. This can in part be found in the technical contributions and solutions provided by the present technology, the utilization of specific training input that is external to a computer, and the providing of the solution or answer external to a computer.

[00150] This section first provides an overview of the current online advertising technology and monetization methods in order to establish a baseline for comparison with novel and useful systems and methods of the present technology (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 (or intelligent entity) attention (4.3) and challenges to the current online Ad monetization model (4.3a) are discussed.

[00151] Since the present technology seeks to improve on the existing online Ad monetization model by overcoming limits to AI systems, these limits are also elucidated (4.4), with specific emphasis on the data challenges (4.4a) facing efforts to develop advanced AI and AGI systems. To show how the inventive system and methods for online advertising can power AGI, the previously disclosed inventive approach to AGI is summarized (4.5).

4.1 Overview of Online Advertising and Monetization

[00152] The present existing technology for online advertising works generally includes the following steps (as further described in FIG. 13):

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 profile and other information (e.g. cookies, user preferences, information about user's past online behavior) to help its Ad targeting system target users to display the Ad to.

3. AI systems use information in (1), (2), and other information to determine the optimal locations, frequencies, times, and users to display the Ad to.
 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, with the hope that the advertiser derives some benefit in terms of marketing or selling their products or services that are reflected in the Ad content.
- 10 **[00153]** Note that traditionally user attention has referred only to human user attention. However, the current inventive technology, system, and methods, recognize that the attention of all intelligent entities – both human and non-human – is valuable. Therefore, the current technology for harnessing and monetizing human attention, and for using that attention and expertise to power AI/AGI/SI systems can be applied equally to the attention of any intelligent entity.
- 15 **[00154]** The basic business model for monetizing user attention is very old. It is an advertising model, similar generally to the advertising models that have been in existence for hundreds of years. Earlier, less technologically advanced versions of the existing advertising model have been used for newspaper, televisions, radio, and other media advertising for many years. The basic idea has been to capture a few seconds of human attention and hope to monetize that attention by converting it
- 20 into purchase behavior by the user.
- [00155]** At a high level, humans are considered to be consumers and their attention is valuable only to the degree that it can be used to 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
- 25 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 purchase behavior of the consumer, even if intermediate steps might be to monetize personal user information by selling to third parties that ultimately try to secure the purchase behavior.
- 30 **4.2 Current Use of AI and Online Ad Targeting**
- [00156]** Online advertisers leverage technology, AI, user data, and customized targeting, display, monitoring, and delivery technology. AI is used as a means for increasing the effectiveness of the Ad targeting and display to increase the CTR or conversion rate and thus justify a higher price per click or per Ad impression. The more data that 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

- 5 [00157] A major 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 to offer premium, subscription-based, versions of technology products that are ad-free. For example, YouTube® offers a version of its service without Ads but requires users to pay a monthly subscription fee in order to view ad-free content.
- 10 Many other online services have similar offerings despite the fact that users opting for an ad-free experience are often those who use the service most heavily and therefore represent the largest potential ad-revenue segment of users.

- [00158] A second challenge comes in the form of ad-blocking technology that eliminates pop-up ads and generally attempts to provide an ad-free user experience without any compensation to the
- 15 content provider.

[00159] A third challenge comes in the form of increasingly restrictive privacy laws and regulations that attempt to limit the use of user data that is used by AI to target ads.

- [00160] All three challenges, reflect a fundamental conflict between the user's interest to view content or use online services unhindered by advertising and the advertisers' need to interrupt the
- 20 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 to do so.

4.3 Current Monetization of Human Attention

- 25 [00161] 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.

- [00162] 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: $2000 \text{ seconds} / 3600 \text{ seconds per hour} = .5555 \text{ hours}$. \$3.12 paid for .5555 hours of human attention equates to an hourly rate of $(\$3.12 / .5555 =) \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/hr. in states like California. Since every business in the US is legally required
- 30

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

5 [00163] 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.

[00164] Unless the fundamental business model underlying online ads were to change, advertisers
10 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.

[00165] 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
15 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.

[00166] 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
20 in emails or in conversations overheard by smart speakers. The backlash has been increasingly intense privacy regulations. 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 re-thinking the fundamental online model itself.

25

4.4 Current Limitations on AI / AGI Development

[00167] Just as online advertising faces limits on monetization, AI and AGI development also faces 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
30 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.

[00168] For example, OpenAI's LLM, GPT 3, 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 is more than 8X increase in about 18 months. Moreover, estimates are

that LLM parameter-counts are now increasing at the rate of about 10X per year. Since, as a first order approximation, 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.

- 5 **[00169]** Moore law predicted a doubling of computational power every 18 months. Requiring 10X more power every year cannot be met by improving chips alone. More chips are required. Given 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
- 10 to meet the exponentially increasing demand for computation power to train the next generation of models. Briefly, due to computational constraints alone, conventional AI and AGI development is facing a severe bottleneck.

4.4a Data Challenges Facing AI / AGI Development

- 15 **[00170]** 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
- 20 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.

- [00171]** Although each of us humans like 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
- 25 average. And it should come as no surprise that the models trained with such data behave in average ways, making many mistakes.

- [00172]** 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
- 30 greater constraint on AI development than obtaining the computing power needed to train models in the existing paradigm.

[00173] 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. The present technology 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.

5

4.5 Description of Previous Inventive Approach to AGI

[00174] 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
10 limitations described above. FIGS. 1 – 13 describe some of the major components of this novel approach to AGI development. FIGS. 14 – 20 describe additional, completely new, inventive components described in this disclosure.

[00175] One reason AGI has been so elusive is that specific knowledge and expertise from diverse fields must be creatively combined in an invention to achieve AGI. Another reason the development
15 of AGI has been non-obvious, is that almost all AI researchers are focused on trying to improve existing narrow AI systems via ever more complex and extensive machine learning approaches.

[00176] The fact that AGI has resisted attempts by thousands of others -- despite the expenditures of huge sums of money -- and the fact that specialized knowledge in relatively obscure fields had to be combined with mainstream AI approaches in the present technology, argue strongly for the novelty
20 and creativeness of the present technology.

[00177] The present technology describes the system and methods not only to achieve AGI, but also to achieve it rapidly, and most importantly, safely.

[00178] It is possible to influence the evolution of AGI in a positive direction. The best way we can do this is by adopting the safest possible path to the development of AGI and ensuring that humanity
25 follows that path. In turn, the best way to ensure that humanity follows the safest path, is to show that the safest path to AGI is also the fastest and therefore most desirable path to AGI. These considerations, the desire to illuminate the fastest path, which is also the safest path, is motivation for the development of the present technology.

[00179] While the above-described devices fulfill their respective, particular objectives and
30 requirements, the aforementioned devices or systems do not describe a system and methods for safe, scalable, artificial general intelligence that allows scaling by using a combination of human users and multiple AI systems to train other AI systems by combining values and ethical knowledge of the human users and the multiple AI systems for training. The present technology additionally overcomes one or more of the disadvantages associated with the prior art.

[00180] A need exists for a new and novel system and methods for safe, scalable, artificial general intelligence that can be used for scaling by using a combination of human users and multiple AI systems to train other AI systems by combining values and ethical knowledge of the human users and the multiple AI systems for training. In this regard, the present technology substantially fulfills this need. In this respect, the system and methods for safe, scalable, artificial general intelligence according to the present technology substantially departs from the conventional concepts and designs of the prior art, and in doing so provides an apparatus primarily developed for the purpose of scaling by using a combination of human users and multiple AI systems to train other AI systems by combining values and ethical knowledge of the human users and the multiple AI systems for training.

[00181] In the following description, for purposes of explanation and not limitation, specific details are set forth, such as particular embodiments, procedures, techniques, etc. in order to provide a thorough understanding of the present technology. However, it will be apparent to one skilled in the art that the present technology may be practiced in other embodiments that depart from these specific details.

[00182] It can be appreciated that the present technology provides a technical effect, contribution and solution with a technical implementation of multiple customized AAAI systems communicating over a collective intelligence neural network, in combination with all the AAAI systems each utilizing a common cognitive architecture including one or more problem solving protocols for generating one or more solutions or answers to a problem request, and providing the solutions or answers to a user for approval. Where the customization of the AI system resulting in the AAAI includes input from human users for training the AI or the AAAI. Further technical contribution or solution can be where the multiple customized AAAI systems can include one or more cloned AAAIs that can each be customized independently of a parent AAAI and independent of other cloned AAAIs of the same system.

[00183] Still another technical contribution and solution is for the faster and safer creating of scalable AGI that utilizes human input in training and customization for imparting human ethical attributes to the AAAI and/or AGI.

[00184] Still yet another technical contribution and solution is for scalably train AI systems and/or agents with a combination of safety and ethical information from many individual AI agents to achieve a representative and statistically valid sample of human ethics and values covering a wide range of scenarios. A further technical contribution can be found in that the present technology includes methods for combining the information from many agents and assembling optimal combinations of such agents for providing scalable training of AI or AGI.

[00185] It can be appreciated that the present technology is found outside of computer program exclusion and/or abstract idea interpretation. This can in part be found in the technical contributions and solutions provided by the present technology, the utilization of specific training input that is external to a computer, and the providing of the solution or answer external to a computer.

- 5 [00186] The AAAI approach to developing safe AGI is fundamentally a Collective Intelligence (CI) approach. The source of intelligence is not a monolithic LLM, SLM or super-advanced AI, but rather a collection of intelligent agents which can be both human and AI. Component sub-tasks in developing AGI include, without limitation, training individual AI agents, combining knowledge (including without limitation subjective values and ethical knowledge) from different agents
10 effectively and efficiently, scaling the AGI, and continuously improving/updating the AGI.

- [00187] Current approaches – such as RLHF and Constitutional Learning – are failing to effectively and scalably train AI to be ethical and safe. The present technology describes a scalable system and methods that are superior to current approaches. In one aspect, the present technology can include the combination of safety and ethical information from many individual AI agents to achieve a
15 representative and statistically valid sample of human ethics and values covering a wide range of scenarios. The present technology can include methods for efficiently covering a wide range of ethical situations and dynamically addressing new situations as they emerge. Methods for combining the information from many agents and assembling optimal combinations of such agents are also presented. These methods can be used not only to improve safety using ethical knowledge
20 but also to create superintelligent systems that combine many other types of knowledge. Safe AGI and SuperIntelligence can be achieved via the collective intelligence approach described in this description of the present technology. A detailed scenario, using the company META® as an example, illustrates one preferred implementation of the present technology.

- [00188] Methods for dynamically updating knowledge are also presented. Successful
25 implementation of the present technology will increase the chances that AI, AGI, and SuperIntelligence remain aligned with human values even when such systems greatly exceed humans in intelligence.

- [00189] Advanced Autonomous Artificial Intelligence (AAAI) is a set of systems and methods for developing Artificial General Intelligence and SuperIntelligent Artificial General Intelligence
30 (collectively “AGI”) in a rapid and safe manner for the benefit of humankind. In contrast to other approaches to the development of AGI, the AAAI present technology achieves a faster and safer path to AGI by relying, at least initially, on the involvement of (ideally many millions of) humans minds in the AGI training, operation, and safety/supervisory functions.

[00190] Advanced Autonomous Artificial Intelligence (AAAI) is a set of systems and methods for developing Artificial General Intelligence and SuperIntelligent Artificial General Intelligence (collectively “AGI”) in a rapid and safe manner for the benefit of humankind. In contrast to other approaches to the development of AGI, the AAAI present technology achieves a faster and safer path to AGI by relying, at least initially, on the involvement of (ideally many millions of) humans minds in the AGI training, operation, and safety/supervisory functions.

[00191] The AAAI present technology can achieve AGI by enabling users to first customize and clone their own AIs. These customized AIs (AAAI) participate in problem solving and other intellectual activities on a network consisting of other AAAs and humans. Although each AAAI on its own may lack the breadth of skills and knowledge to be an AGI, collectively the AAAs (initially with help from humans on the network) form an AGI that will quickly surpass average human ability in all intellectual endeavors.

[00192] Some aspects of the present technology can include: 1) the system and methods to customize AIs with the unique knowledge, skills, and ethical values of the users; 2) the universal problem solving architecture that allows AAAs to interact productively with each other and with humans on intellectual tasks; 3) the network where the interactions takes place; 4) the methods for integrating the knowledge and ethics of individual AAAs into an AGI; and 5) the methods for learning and continuous improvement so that the AAAs and the AGI become smarter and more ethical over time. Involvement of humans as customizers of their AAAs and participants on the network is an essential feature of the present technology which not only accelerates the development of AGI, but also makes AGI safer by providing a mechanism for the ethical values of millions of humans to be adopted by and reflected in the AGI.

[00193] One implementation of the AAAI system of the present technology has a focus on safety and is implemented via five sub-systems and associated methods, as illustrated in FIG. 1. The five sub-systems of the AAAI system are: 1) AAAI Customization, 2) AAAI Architecture, 3) AAAI Network, 4) AAAI Integration, 5) AAAI Improvement. The acronym SCAN--II (Safe, Customizable, Architecture and Network -- Integrated and Improving) describes the present technology in the exemplary implementation. Other combinations of subsystems, and variations of each subsystem, are also possible. Safety features have been designed into each sub-system in an effort to provide redundant safety checks in the event one or more sub-systems are omitted from a particular implementation.

[00194] The five sub-systems of the AAAI system can be further described as:

- 1) A base level Large Language Model (LLM), Small Language Model (SML), or other AI system can be customized to reflect the knowledge of an individual, group of individuals, or

organization and designated an Advanced Autonomous Artificial Intelligence (AAAI).

- 2) The customized AAAI can be enabled to participate in problem solving using a universal problem solving architecture that is compatible with both human and AI agents.
- 3) The problem solving-enabled AAAI participates in problem solving activity, including but not limited to:
 - 5 planning, problem solving, and other types of sequential, multi-step cognitive activity. on a network of intelligent agents;
 - generate and select operators that reduce a difference between a current state of problem solving and a desired state based on the goal/subgoal;
 - 10 setting of a subgoal towards achieving the goal;
 - utilizing hierarchy until an actionable goal is set that can be acted on by the operator; and
 - analyzing the auditable record to determine recommendations for improvement of the problem solving process to achieve a solution to the goal/subgoal.
- 4) Multiple AAAIs, or PSIs. on the network can be integrated to achieve AGI; or AI capable of intelligent (or super-human level) behavior across a wide range of tasks.
- 5) The individual AAAIs, the problem solving network, and/or the integrated system of multiple AAAIs continuously improve via a variety of means, including but not limited to, redirecting the efforts of individual AAAIs and/or the integrated AGI towards the task of improving the system and/or components of the system.
- 20 **[00195]** The sub-systems or new sub-systems can include any one of or any combination of:
 - 1) Safety/ethics check - Comparing a goal or subgoal against a list of prohibited attributes and assigning an ethics value based on a result of the comparison. Checking the goal/subgoal against a list of prohibited attributes. Combining values/safety information from AAAIs, using a set of approved criteria for a task by a user or by a regulatory agency or by AAAIs approved by human user. Establishing or using a threshold for the goal/subgoal to determine if the ethics value is unsafe, unethical, safe, or ethical. Determining if a sequence of individually safe goals/subgoals are unsafe or unethical when considered cumulatively. Determining whether a violation occurred reflects a predictive evaluation if the goal is to violate the ethical criteria. Recording any and all activity of the safety/ethics check in the auditable record.
 - 25 2) AAAI matching - Detecting and identifying additional AAAIs that each have a criteria related to one or more goal or subgoal criteria.
 - 3) Remembering and/or improving – Recording activity, comparing with successful or unsuccessful progress towards the problem solutions, determining which activity to keep active or forget.
- 30

- 4) AAAI learning – Learning, including a procedural learning process that utilizes information provided by intelligent entities such as human users equipped with computers or AAAIs. Recording activity, comparing with successful or unsuccessful progress towards the problem solutions, determining which activity to keep active or forget. Assigning credit value or blame value to a group of content of the problem solving activity. A set of prompts provided to the user and information received based on the prompts. Updating AAAIs with the group of content determined as active. The group of content can be, but not limited to, a set of prompts provided to the user and information received based on the prompts, all of which being recorded in the auditable record. Optionally, the problem solving activities can include the group of content.

10

EXAMPLE USER SCENARIOS

[00196] It may be helpful to describe some user scenarios that provide a sense of how the present technology can operate in some of the aspect implementations. An exemplary process is illustrated in FIGS. 2 & 3.

- 15 [00197] In one aspect, a user “visits” AAAI.com via the user’s computer, cell phone, PDA, or goggles. AAAI.com would interact with the user via a web-based interface, a phone app, custom software for the PDA, or a metaverse/ virtual reality environment. The mode of interaction could be physical via a keyboard, mouse, or gestural interface; voice-based via a microphone input coupled to natural language understanding and generation systems; or video-based as in the case where the user becomes an avatar in a virtual reality setting or in the metaverse.

- 20 [00198] The initial interaction would include setting up the user’s account, which might be free or paid. This would involve an account name and password or other authentication mechanisms which might include, without limitation, biometric forms of ID such as fingerprint, face or voice recognition, and/or multi-factor authentication mechanisms such as software or hardware authenticators residing on a separate security device or on one of the user’s existing devices.

- 25 [00199] For security, all communication between the user and the AAAI system could be encrypted via a VPN and/or could use other methods of encryption and security which are well known in the art of programming.

- 30 [00200] AAAI.com may request that the user set up payment capabilities via credit card, PayPal, Venmo, blockchain, ACH, or other payment mechanisms. These payment capabilities would allow funds, payments, and/or credits to be transmitted bi-directionally – from the user to the AAAI.com and also from the AAAI system to the user in cases where the AAAI system needs to pay or credit users for work efforts of their AAAIs or broker payments between users and/or between AAAIs on the AAAI network.

[00201] In one aspect of implementation, AAAI.com can have interfaces with other companies and vendors that the user might use -- including, without limitation, and for example: Facebook, Instagram, Reels, Amazon, Apple, Microsoft, Google, and YouTube.

5 [00202] In the initial interaction with the user, and subsequently upon user request, AAAI.com would engage in a dialog or other interaction (which could include presenting the user with menu options, lists, graphics, sliders, buttons, and other user interface controls in a GUI, textual, haptic, voice, or VR-related manner) with the user to determine the user's goals and objectives in using the AAAI system.

10 [00203] For example, some of the objectives a user may have in using AAAI.com may include creating and customizing their own AI (known as an AAAI) for purposes that might include, without limitation:

Serving the user as an advisor, teacher, or companion.

Representing the user in negotiations, interactions, discussion, and transactions with other users, or with the AAAIs of other users; or with vendors and other companies.

15 Working on behalf of the user for compensation, or in volunteer efforts, where such work includes online intellectual, advising, or problem-solving work across a wide range of tasks.

Duplicating or "cloning" the user's AAAI so that several or many of the cloned AAAIs can work on behalf of the user in parallel, including interacting with, teaching, and improving each other so that the cloned AAAIs increase their knowledge, skills, and abilities.

20 Serving as legacy AAAIs that can continue to interact with the world, including potentially comforting living relatives and friends, after the owner's death.

Contributing knowledge, ethics, and effort to AAAI.com's AGI, and improving the base level of AI or AGI that AAAI.com can offer users before those users add their unique customizations.

25 Working with other users' AAAI to help ensure ethical and safe behavior by AGI by contributing ethical information and values to the AGI and participating in monitoring, review, supervision, and voting processes that can help ensure the AGI remains safe and ethical.

30 [00204] In the dialog or interaction with the user, the AAAI system will also identify constraints and resources available for customizing the user's AAAI. For example, some of these constraints and resources, might include, without limitation:

The amount of training and/or supervisory time that the user has to devote to customizing their AAAI.

The amount of financial resources the user is willing devote to customizing their AAAI.

Availability of social media information such as Facebook profiles and timelines, Instagram profiles and histories, Reels, TikTok, and YouTube videos, tweet and text content and histories, emails and email histories, cookies collected by advertisers, blog posts, articles, books, patents, audio and video recordings, pictures, and other information about, and/or collected by, the user or third parties that could be used to train, tune, or customize the user's AAAI.

Availability and use of personality tests, such as the Myers-Briggs personality inventory, skills and knowledge assessments, standardized tests, exams, certifications, and other types of assessments and questionnaires which could be given online (or which have already been given) to the user.

Availability and use of other knowledge bases and training data from users on the AAAI platform that could be used to train, tune, or customize the user's AAAI.

Other human users, and/or their AAAs, available to help train, tune, or customize the user's AAAI.

Other texts and information, individual texts, and libraries selected by the user or by the system for purposes of training the user's AAAI. For example, the Bible, Koran, Dhammpada, Mahabharata, or other spiritual/ethical/religious texts might be selected for training the AAAI based on the user's religious preferences; books on plumbing might be selected if the AAAI will be used to primarily solve online plumbing problems. Even if these materials are part of the base AAAI that is provided to the user, emphasizing certain texts or subsets of information for additional training can result in the user's AAAI's behavior being more reflective of how a plumber, or Muslim, or Christian might behave, for example.

[00205] In addition to specifying objectives, resources, and constraints via an interactive dialog or other interaction with the system, the user or system may want to specify other technical parameters that affect the training or customization process. These parameters can include, without limitation:

The type of training, tuning, or other ML algorithms that are used.

The type and size of the training dataset(s).

The degree to which the training materials are to be "cleaned", formatted, labelled, or otherwise processed before customization begins.

The number of training "epochs" or iterations through the learning algorithm(s).

The sophistication and type of base model(s) being customized or trained.

The required timeframe for training – e.g., must be completed in a minute, a day, a week – which might have implications for cost and resources used.

The “temperature” or other parameters internal and specific to various machine learning algorithms that can affect what is learned and how it is learned including, without limitation, how literal or how divergent or “creative” the customized AAAI will be in its responses.

Whether “one shot”, “few shot”, or extensive training is to be used.

5 The amount of human and/or AI supervision to be used in the customization process.

[00206] Once the user’s AAAI is customized, the user can clone it and/or put it to work on the user’s behalf on the online network. The user’s AAAI can begin acting on the user’s behalf making travel arrangements (for example), providing advice, interacting with other AAAIs, participating in the collective AGI efforts by contributing problem solving as well as ethical information, and
10 potentially earning money on behalf of the human user.

SIMPLE EXEMPLARY IMPLEMENTATION

[00207] FIG. 3 shows one simple exemplary implementation of the system and methods for creating an ethical and safe Artificial General Intelligence from the collective intelligence of AAAIs and
15 humans. This simple implementation is compatible with all of the company and platform specific scenarios outlined above, as well as with many other potential integration scenarios. FIG. 3 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.

[00208] A (human, AAAI, or other intelligent entity) user visits the AAAI.com website (a). The
20 website informs users and offers them two actions: Sign Up (b) or Login (c).

[00209] If the user opts to Sign Up then a dialog is initiated which 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 money budget.

[00210] 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
25 knowledge (l). They also have the opportunity of buying training from other AAAIs on the network (m).

[00211] After making time (and optionally money budget (h, l, 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
30 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 (o) the user directs 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.

[00212] The AAAI now begins to learn (p). There are two main ways of learning, automatic (q) and human (r).

- 5 [00213] Automatic learning includes, without limitation, learning by interacting with copies of itself (s), learning via interactions with other (optionally supervised) AAAIs (t).

[00214] Human learning includes interaction with humans, either the owner (u) or other humans on the network (v).

- [00215] Both humans and AAAIs can supervise learning of an AAAI. After each (automatic or
10 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.

[00216] At any time, the user can purchase additional training modules (h) that have been proven to increase an AAAIs abilities.

- 15 [00217] The human sets a performance criteria (w) after which the AAAI goes LIVE (x).

[00218] Once live, the AAAI can visit the WorldThink Tree (y) and Browse (z).

[00219] The AAAI can enter the tree as either a worker (a1) or a client (b1).

- [00220] Workers are automatically matched (c1) to tasks or they can select a specific task via search
20 (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).

[00221] Clients (b1) can specify objectives (k1) which are combined with the values/ethics (d), and prior goals and objectives (e) for the system to solve.

- [00222] The client can request that only his/her/their AAAI be used in which case problem-solving
25 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).

[00223] The system can also place non-profit humanitarian and ecologically-oriented tasks, as well as tasks that are part of Planetary Intelligence, on the WorldThink Tree (m1).

- 30 [00224] Clients might (optionally) authorize the system to use copies of their AAAI and data for these purposes without remuneration in exchange for maintaining and operating the free AAAI network when they created their AAAI (n).

Additional Comments on Exemplary Implementation Shown in FIG. 3

[00225] We now provide additional comments on the various elements of FIG. 3, including without limitation, some potential integration points with the illustrative partners mentioned above.

[00226] The “website” (a) could be hosted on Amazon AWS, Microsoft Azure, Google Cloud, Apple Cloud, Nvidia datacenter offerings – or could have native implementation on the platforms of any large tech company. “website” could also be an “app” in the AppStore or other App marketplace. It could be a government-sponsored, nonprofit, or other globally-accessible technology that is able, directly or indirectly, to link some of the attention of all human beings who wish to participate. Also, browser plug-ins could be used whereby AAAs 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 AAAs with these files. The “website” could also be an API or other means for connecting AAAs or non-human intelligent entities directly to the network.

[00227] Sign Up or

[00228] Login (c)- could be via Facebook, Instagram, Apple, Microsoft, Google, You Tube, Tik Tok, Amazon, or any other partner ID scheme. Multi-factor authentication and all best ID and security practices enabled. In the event of a browser plug-ins or apps, login to these technologies could serve as login to the AAAI account.

[00229] Values and ethics (d) are elicited via a series of scenarios that have been customized for the user and that are 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) in order 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).

[00230] Goals and objectives (g), together with the budget of time and/or money allocated to reach objectives are elicited via a series of dialogs and/or custom interactions with the system. Budget refers to 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 are helpful in determining 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. Data from partners, reflecting user preferences and other user behavioral information, could also be used by the system to help infer or deduce user goals and objectives.

[00231] 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 supervising the

AAAI, users can ensure that their AAAs meet client goals and expectations – especially in areas where the AAAs 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 by, without limitation, determining goals and sub goals, are ways that human users can assist their AAAs in problem-solving. Generally, by providing human expertise in areas where AAAs are not as proficient as humans, overall problem-solving and the overall effectiveness of the AGI network is increased.

[00232] (g, II) “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 (j) is indicated as debiting the client account (II), of course the worker’s account would also be credited. Generally, a user’s account can be viewed as both a client account and 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 AAAs 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 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.

[00233] (h, i, j, k, I) Training modules (h) could be offered by AAAI.com or by third party partners, including, without limitation, any of the 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) (I).

[00234] (m) AAAI knowledge is a specific type of knowledge that has been already learned by other AAAs, and which can be transferred to a new user AAAI. Such knowledge may not be packaged in the form of a module (e.g., module on accounting) but rather as 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 AAAs with a reputation of 5 stars or higher in the domain of plumbing”.

[00235] (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 AAAs to create AGI-level AI.

5 [00236] (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 given permission for 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
10 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 a 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 AAAs.

15 [00237] (p) Training refers to the process whereby the AAAI is trained or tuned on data, including feedback from the user, other humans, and/or AAAs (including, without limitation, copies of, and variants of, itself).

[00238] (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 as well as with (optionally) other AAAs in order to improve via the interactions. If humans are sometimes involved in the training loop (t) that 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 (s). Just a chess AI can quickly evolve from novice to Grandmaster ability by simulating millions of chess
25 games very quickly, an AAAI can quickly evolve its abilities by simulating many 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.

[00239]

[00240] Humans (or AAAs) can specifically target types of scenarios for automatic learning so that
30 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 AAAs 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 AAAs “on demand” with the skill sets that are needed at any particular point in time.

[00241] 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 AAAs 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.

[00242] (**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 his/her/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 standard before allowing AAAs to work on their platforms, products, applications, or networks.

[00243] (**y, z, a1, b1**) The WorldThink Tree 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 AAAs 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: Worker (**a1**) or Client(**b1**). 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.

[00244] (**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 AAAs. Workers can be matched and recruited from partners (e.g., LinkedIn, Mechanical Turk, Facebook) that have data on human users and/or their AAAs. Workers can also be recruited via online ads offering work on various tasks and targeted to potential workers using ad-targeting mechanism that are well known in the art or described in other patents by the applicant.

[00245] (**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 for AAAI workers).

[00246] (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 as clients that
 5 build upon existing problem solving work.

[00247] (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).

10 [00248] (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 AAAs 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 Planetwide basis (aka “Planetary Intelligence,”).
 15 Various partner organizations – including non-profits, governments, and charitable organizations – might “plug in” their tasks, problems, goals, and objectives here (m1).

[00249] (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
 20 authored by the applicant, this and other PPAs related to AAAI, together with modifications and variations to reflect different modes of reward, payment, and operation.

[00250] 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
 25 and other sites and the data can populate the WorldThink Tree to increase its comprehensiveness.

[00251] 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
 30 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.

[00252] (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.

[00253] 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 re-used when similar problems arise in the future.

[00254] Optionally, royalties may be enabled so that if a user’s or the user’s AAAI’s solution is re-used, a fee is paid to that user in the form of a royalty on the solution. Such royalties can (optionally) be made using “smart contract” on the blockchain or via other payment methods.

[00255] (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.

[00256] The WorldThink protocol is a problem-solving architecture that can be used by AAAI.com to serve as a universal problem solving architecture as it incorporates the general architecture of HPS while adding features to overcome certain challenges.

[00257] In some embodiments and as generally illustrated in FIG. 5, the procedural learning process can occur within the common cognitive architecture.

[00258] The shared and universal problem solving architecture, illustrated in FIGS. 4, 6, and 10 can be exemplified by the following scenario, mentioning humans but also applicable generally to any intelligent entities.

- 1) Problem descriptions can be entered into the AAAI.
- 2) Then human (or intelligent entity) problem solvers can be identified and recruited into a database or data source of human workers.
- 3) Qualified humans or intelligent entities can be matched to problems.
- 4) Use LLMs or other means to translate English descriptions of problem tasks, goals, operators, and solution steps into language of a universal problem solving architecture.

- 5) Delegate work on sub-problems to different human (or intelligent entity) problem solver(s) so that work on multiple aspects of a complex problem can proceed in parallel.
- 6) Combine solutions to various sub-problems into an overall solution.
- 7) Direct the attention of problem solvers to parts of the problem tree where their work is needed.
- 5 8) Compensate or pay workers for solutions to the problem and/or sub-problem(s).
- 9) Allowing human user to accept the solution, reject the solution, and/or provide feedback to solvers on their solutions to the problem and/or sub-problem(s).

[00259] Referring to FIG. 5, the steps of solution learning can be exemplified with the recording at each step of the learning process operators applied, new state of the problem, evaluation function
10 used and its results, current relevant goal/subgoals, and other information that differs from previous step(s). The state of the problem or problem state can be evaluated to determine if the problem is solved. If not, then using information from the latest problem state after the last step, re-run the problem-solving process, evaluation of progress, and selection of next operators to apply. After which, the process can return to the step of recording.

15 [00260] If the problem is solved, then record successful or unsuccessful solutions for retrieval to save effort of solving previously solved problems and to inform problems solving efforts about previous unsuccessful paths.

[00261] Successful solutions and unsuccessful attempts with key words for future matching/retrieval can be indexed using semantic analysis, hash functions, and/or other means.

20 [00262] A periodical review of all stored solutions can be implemented to ensure they meet established ethical and safety guidelines, and flag unsafe/unethical solutions for removal from the database or data source.

[00263] Periodically update and propagate changes to the solution database so problem-solving network and agents can access an ever-increasing repertoire of solutions as well as increasing
25 knowledge of unsuccessful attempts.

[00264] Referring to FIGS. 6 & 7, the present technology can include a utilization of a network of multiple intelligent entities including human workers in combination with a universal problem solving architecture. The multiple intelligent entities are matched to a problem request based on a problem criteria using a database or data source including a list of human and/or AI problem
30 solvers. Any part of the problem request can be translated into an unambiguous language utilizing a universal problem solving architecture including the decision tree.

[00265] A sub-problem of the problem request can be delegated to one or more of the matched intelligent entities so that work on the sub-problem proceeds independently from each other and parallel with each other, as further illustrated in FIG. 12. The universal problem solving architecture

is utilized in a problem solving process on the sub-problems, respectively, to create one or more sub-solutions.

[00266] Any one of or any combination of the intelligent entities can provide in natural language a description of any one of or any combination of a current problem state, a goal of the problem request, relevant problem solving information, and a next step that the human (or intelligent entity) workers will take in the problem solving process.

[00267] The sub-solutions can be received from each of the matched intelligent entities for the sub-problem delegated thereto. Any one of or any combination of the sub-solutions and an overall solution can be provided to any one of or any combination of a user interface of a user AI system or the intelligent entities.

[00268] Parsing and translating, by the intelligent entities, the natural language description into the unambiguous language can be utilized by the decision tree of the universal problem solving architecture.

[00269] In some embodiments, if the intelligent entities are unable to specify a problem state, including relevant operators and information needed to take a next step in the problem solving process based on the parsing and the translation, then the intelligent entities can engage in dialog with at least one of the human workers until a precise problem state is specified.

[00270] Some embodiments the problem solving process can be repeated until the overall solution is accepted or resources are exhausted. The matched human (or intelligent entity) workers can be compensated for the sub-solutions, respectively. Further, a reputation attribute can be assigned to any one of or any combination of the human (or intelligent entity) workers and the worker AI system.

[00271] In some embodiments, the solving process can include a series of problem state transitions from an initial problem state where there is a goal to a final solution state where the goal has been achieved, and wherein a series of decisions are made by the problem solving process and actions taken that applies operators that enable the human workers to transition from state to state until the final solution state is reached.

[00272] Referring to FIGS. 4 & 10, the present technology can include a utilization of a network of human (or intelligent entity) users in combination with a universal problem solving architecture. The multiple human (or intelligent entity) users are matched to a problem request based on a problem criteria using a database or data source including a list of human and/or AI problem solvers.

[00273] A sub-problem of the problem request can be delegated to one or more of the matched intelligent entities so that work on the sub-problem proceeds independently from each other and parallel with each other, as further illustrated in FIG. 12. The universal problem solving architecture

is utilized in a problem solving process on the sub-problems, respectively, to create one or more sub-solutions.

[00274] The sub-solutions from each of the matched human (or intelligent entity) workers can be provided for the sub-problems delegated thereto. The matched human (or intelligent entity) workers
5 for the sub-solutions can be compensated, respectively.

[00275] Any one of or any combination of the sub-solutions and an overall solution can then be provided to a user interface of a user AI system or any other AI system.

[00276] The human (or intelligent entity) user is allowed to accept the overall solution, reject the overall solution, and/or provide feedback to any one of the matched human (or intelligent entity)
10 workers on any one of the sub-solutions.

[00277] A reputation attribute can be assigned to the human (or intelligent entity) workers and/or the worker AI system. The reputation attribute can include metrics on any one of or any combination of a time to the sub-solutions, a difficulty value of the problem request, short and long-term user satisfaction with the sub-solutions, a number of times any one of the sub-solutions has
15 been re-used on the network, a rating other human (or intelligent entity) workers, a responsiveness value of the human (or intelligent entity) workers, and a reliability value of the human (or intelligent entity) workers.

[00278] Some embodiments can include using the reputation attribute in the matching of the human (or intelligent entity) workers to the problem request using an algorithm to the delegation of the sub-problems, and/or compensating the matched (or intelligent entity) human workers for the sub-solutions, respectively.
20

[00279] In some embodiments, the algorithm can use a hierarchy of the metrics that is preset by a human user of the problem request.

[00280] Some embodiments can include recording information on each step of the problem solving process by the human (or intelligent entity) workers or the worker AI system.
25

[00281] Some embodiments can include recording a criteria of the recorded step of the problem solving process, the criteria being a time taken for each step.

[00282] Some embodiments can include analyzing the recorded information after the overall solution is accepted or after the problem solving process is complete, and updating the metrics of the reputation attribute.
30

[00283] Some embodiments can include soliciting, at predetermined intervals after the overall solution or the sub-solutions are provided to the user interface, a survey for user satisfaction information to obtain short and long-term satisfaction metrics that are used to update the reputation attribute of one or more of the human (or intelligent entity) workers or the worker AI system.

5 [00284] Referring to FIG. 8, the present technology can include a utilization of human users and AI systems, which includes an execution of a scalable safety/ethics check on any one of or any combination of a goal, and a solution for the goal provided by any one of or any combination of the intelligent entities including any one of or combination of human users each using a computer system and AI systems.

[00285] The goal and/or the solution can be compared against prohibited attributes, and an ethics value can be assigned to the goal and/or the solution based a result of the comparison and/or an ethics criteria.

10 [00286] Based on the result of the comparison, a common cognitive architecture including one or more problem solving protocols can be conducted on the goal to create the solution and thereby creating an AGI. The results of the comparison and the solution can be provided to any one of the intelligent entities.

15 [00287] In some embodiments, the ethics check can be performed at any one of or any combination of when the goal is provided, and periodically from when the goal is provided to when the solution is provided.

[00288] In some embodiments, the ethics criteria can be determined by any one of or any combination of combining values and safety information from one or more of the intelligent entities, using a set of approved ethics criteria mandated for a particular task by a user or by a regulatory agency. It can further be provided by any one of the additional intelligent entities and validated or approved by the human user.

[00289] In some embodiments, the ethics criteria can include a confidence level threshold for the goal so that the ethics value is determined as any one of an unsafe goal, an unethical goal, a safe goal, and an ethical goal.

25 [00290] In some embodiments, the confidence level threshold can be further utilized to determine if a sequence of individually safe goals is unsafe or unethical when considered cumulatively.

[00291] In some embodiments, the confidence level threshold can be utilized to determine whether a violation occurred that reflects a predictive evaluation if the goal is to violate the ethics criteria.

[00292] In some embodiments, a candidate goal can be proposed based on the ethics value, and the candidate goal is compared against the prohibited attributes.

30 [00293] In some embodiments, the results of the comparison can be recorded in an auditable record for use in the determining which problem solving activity leads to the solution to keep active.

[00294] Further referring to FIG. 8, the scalable ethics check can compare any one of or any combination of the problem request, the sub-problem and the sub-solutions against prohibited

attributes and assign an ethics value based on any one of or any combination of a result of the comparison, and an ethics criteria.

5 [00295] In some embodiments, the step of the ethics check can be triggered every time the problem request or the any one of the sub-problems is set by the human (or intelligent entity) user, and/or triggered each time compensation is provided to the matched human (or intelligent entity) workers.

[00296] The goal/subgoal can be compared against a list of prohibited attributes. The ethics criteria can be determined by any one of or any combination of combining values and safety information from any one of the AAAs. Combining values/safety information from AAAs, using a set of approved criteria for a task by a user or by a regulatory agency, or by AAAs approved by human
10 user.

[00297] The ethics criteria can include a confidence level threshold for the problem request so that the ethics value is determined as any one of an unsafe goal, an unethical goal, a safe goal, and an ethical goal. The confidence level threshold can be further utilized to determine if a sequence of individually safe goals is unsafe or unethical when considered cumulatively.

15 [00298] In some embodiments, the confidence level threshold can be utilized to determine whether a violation occurred that reflects a predictive evaluation if the goal is to violate the ethics criteria. Any and all activity of the safety/ethics check can be recorded in the auditable record.

[00299] FIGS. 9-13 provides a simple exemplary framework for understanding the WorldThink protocol. In the implementation using the WorldThink protocol, clients pay for solutions using
20 tokens. The solutions are produced by harnessing the collective power of many human (and machine, or AAA) intelligences. Clients can use different domain-specific AAAs for different types of problems.

[00300] The WorldThink protocol is the foundation of the pyramid. The protocol layer provides an (optionally, Ethereum or blockchain based) infrastructure that makes it much easier for developers
25 to build and scale customized problem-solving AAAs. The protocol enables re-use of solutions within and across AAAs. It also handles payment of royalties via smart contracts, reputation metrics, and other functionality that assists AAA customizers and developers and promotes network effects.

[00301] FIG. 9 a diagram illustrating various use cases for domain-specific problems which depend
30 upon the underlying WorldThink protocol, and which together help form the basis for an AAA and/or AGI system capable of solving a wide range of problems. At the top of the pyramid are Collective Intelligence Solutions. Integrating the Collective Intelligence of AAAs (and human problem solving agents) is the means to achieve AGI, as discussed earlier.

[00302] In the implementation using the WorldThink protocol, clients pay for solutions using tokens. The solutions are produced by harnessing the collective power of many human (and machine, or AAAI) intelligences. Clients can use different domain-specific AAAIs for different types of problems.

- 5 [00303] The WorldThink protocol is the foundation of the pyramid. The protocol layer provides an (optionally, Ethereum or blockchain based) infrastructure that makes it much easier for developers to build and scale customized problem solving AAAIs. The protocol enables re-use of solutions within and across AAAIs. It also handles payment of royalties via smart contracts, reputation metrics, and other functionality that assists AAAI customizers and developers and promotes
- 10 network effects.

[00304] In the exemplary, FIG. 10 shows a simple exemplary universal problem solving framework under the common cognitive architecture, and which can include:

- defining a problem space configured or configurable to support all possible states of the problem request, the states including any one of or any combination of an initial state, a goal state,
- 15 and all intermediate states that can be reached from the initial state;
- applying means-ends analysis on the problem request to break the problem request down into goals and subgoals by identifying a difference between the current state and the goal state, and then applying the operators to reduce the difference, a safety or ethics screening is applied each time the goals or the subgoals is set;
- 20 applying heuristic rules that are configured or configurable to guide the selection of the operators in an absence of the completion solution, the heuristic rules are used to reduce the problem space;
- identifying one or more second operators configured or configurable to enact an action to transform one of the states into another state, the second operators move from the initial
- 25 state to the goal state by changing a current state of the problem request;
- applying a control structure including a set of rules that govern a selection of the second operators to be applied at each step of the problem solving protocols, and that determines which of the second operators to apply next based on the current state of the problem request and the goal state;
- 30 applying evaluation functions to determine an application of the second operators;
- assigning a credit or blame value to the completion solution or sub-solution to the completion solution that enables tracing back and determining which of the second operators were most useful and also which of the evaluation functions led to success or failure of problem solving attempts;

recording of both successful and unsuccessful problem request solution attempts; and
analyzing the solution attempts to improve selection of the heuristic rules and the evaluation
functions.

[00305] Existing collective intelligence approaches to problem solving have been largely limited to
5 simple one-step approaches, such as those used by question and answer (Q&A) systems (e.g.,
Quora, Google Answers, Yahoo Answers). LLMs such as GPT also largely fall into the category of
Q&A systems since they were designed to generate responses given an input, rather than to solve
problems per se. While such Q&A systems have had some success at simply aggregating the
responses of many online participants, these systems are not designed to handle complex, branching,
10 multi-step problems. Simple aggregation of responses (or even betting on outcomes as seen in
prediction market approaches such as Augur and Gnosis) is quite different from coordinating the
efforts of many respondents to solve complex problems. The WorldThink protocol is specifically
designed to overcome the challenges inherent in coordinating many intelligent entities to represent
and solve complex, multi-step problems in an automated way that fairly rewards participants.

15 [00306] In the exemplary, FIG. 11 shows some of the basic problem-solving functionality supported
by the WorldThink Protocol, generally referenced with numeral 10.

[00307] Problem solving begins when a client on AAAI.com submits a problem-solving request to
the community of online participants (Step 12). All AAAIs, or human solvers, following the
protocol gather certain standard information from the client. A partial list of this information can
20 include: the name and description of the problem, the total reward that the client will pay for a
successful solution to the problem, the criteria to determine whether a solution will be deemed
successful, the time limit for solving the problem, the minimum and maximum number of problem
solvers allowed to work on the problem simultaneously, qualifications required of participants
working on the problem, which parts (if any) of the problem and solution will be confidential,
25 whether the solution must be exclusive to the client or whether it can be re-used for others, and
parameters relating to how to reward multiple problem solvers for their efforts and/or successful
solutions.

[00308] The client can break complex problems down into a series of sub-problems or request that
the community take on this task as part of the problem-solving effort. The client user-interface,
30 which could be a dialog initiated by an AAAI can be customized by the AAAI owner, but the
underlying data format is standard and specified by the WorldThink or ODPS protocol. Once the
client has submitted a problem, AAAI.com can recruit participants using its own custom methods
and/or leverage recruiting and reputational screening functionality that is built into the WorldThink
protocol and thus shared by all AAAIs.

- [00309] Solvers work on the problem following a rigorous structured problem-solving process that is common to all problem-solving agents and enforced by the WorldThink Protocol (Step 14). For example, each step in the problem-solving process must be in service of a named goal and must take a named action in order to transition the problem solving from the current state to the next state.
- 5 Every problem-solving step is represented in a decision tree which is supported by the protocol (optionally captured in Ethereum logs) and which participants can view via AAAI.com.
- [00310] When a Solver submits a complete solution (Step 16), it is timestamped and validated against the client's success criteria before being passed on to the client (Step 18) for final acceptance. Once the client accepts the solution, smart contracts can automatically distribute tokens
- 10 to the problem solver based upon the problem payment parameters (Step 20) or other, more centralized, payment procedures can be used.

Collaborative Problem Solving Using the WorldThink Protocol

- [00311] In the exemplary, FIG. 12 shows the same steps in an example where two problem solvers
- 15 (which could humans, AAAIs or a combination) collaborate to solve a client problem, as generally referenced with numeral 22. In this case, the overall problem has been broken down to include a sub-problem. Solver 1 has expertise in assembling an overall solution but cooperates with Solver 2, who provides a solution to the sub-problem (Steps 30 and 32). When the overall solution to the problem is submitted to the client (Step 34), rewards are paid to both Solvers (Step 36) based on the
- 20 objective record of their contributions and the agreed upon payment parameters.

- [00312] The WorldThink protocol supports breaking problems into sub-problems in several ways. First, the client may choose to specify sub-problems when submitting the overall problem (Step 24). Alternatively, Solver 1 might begin working on a problem and realize that the total solution requires solving a sub- problem outside of his/her expertise. Solver 1 could then create a sub-problem,
- 25 offering up a share of the problem's total token reward to anyone who helps solve the sub-problem. Solver 2, who has the required expertise and who can see the new sub-problem posted by Solver 1 on the decision tree. The decision tree may be optionally maintained in Ethereum logs, or via a centralized method. The solvers access the tree via AAAI.com (or optionally directly from the blockchain). Then Solver 2 can work on the sub-problem and submit a sub-solution as part of Solver
- 30 1's overall solution.

[00313] There can be many "Solver 1s" working on the client's problem in parallel, each of whom may be posting sub-problems to attract multiple "Solver 2s". Problem solvers (human or AAAIs) are motivated by the rewards and payment rules associated with (sub) problems. They also care about the quality of work done so far (which is timestamped, attributed, and recorded auditably in

Ethereum logs to ensure transparency and fair assignment of credit) as they choose which (sub) problems to work on. Working on quality sub-problems is more likely to lead to token rewards. This market mechanism helps ensure efficient, fair, and cost-effective solutions.

5 [00314] FIG. 13 is a diagram illustrating features and functions of the Problem Solving Tree structure in the WorldThink protocol. A hierarchical tree construct is created that represents all problem-solving activity by the user, AAAI and/or additional AAAIs.

[00315] Data structure can be utilized that is navigable by the AAAI and/or additional AAAIs to access any part of the problem-solving activity on any part of the hierarchical tree construct.

10 [00316] Searching can be performed on the data structure to locate a predetermined reward associated with the goal and/or the subgoal.

[00317] A matching operation can then match AAAIs to problems or subproblems.

5.0 Opportunities and Benefits Enabled by the Present Technology

15 [00318] Now that we have reiterated the novel approach to AGI covered in 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 present technology, given this context.

[00319] First, the present technology provides a way to increase the intelligence of AI/AGI systems by overcoming the data bottleneck that AI researchers currently face.

20 [00320] Second, the present technology provides an opportunity to monetize human (or intelligent entity) attention far better than the existing online advertising model which is suffering from diminishing returns on incremental Ad targeting improvements.

[00321] Third, the present technology helps solve the most important problem related to the development of advanced AI/AGI systems, namely AI safety, including without limitation, the
25 “Alignment Problem.”

5.1 Opportunities to Increase Intelligence of AI / AGI Systems

30 [00322] 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 resource 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 present technology seeks to increase the intelligence of AGI systems by supplementing the computational power of GPU 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 a tremendous expense to perform the same cognitive task(s) at an inferior level. 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. Note that in cases where the attention of non-human intelligent entities is desired, the current technology enables harnessing and monetizing this non-human attention as well.

[00323] 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 that was described above, and illustrated in FIGS. 1 – 13, 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.

[00324] 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 was 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.

[00325] The present technology 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 present technology recognizes that only certain data is needed at any given time. The present technology focuses 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 (or intelligent entities), over time.

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

- 5 [00326] 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. This 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, education of the human whose attention is being tapped.
- 10 [00327] 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
- 15 over the current online Ad model of between 62X – 178X!
- [00328] Actually, the theoretical improvement in monetization of attention enabled by the present technology is even higher than this for two reasons:
- 1) The present technology 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.
 - 20 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

- 25 [00329] 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
- 30 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.
- [00330] 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 pipe. The

whole episode took no longer than 5 minutes, including you asking a couple of questions. How can you charge us \$10,000?"

[00331] "Ahhh," said the plumber with wink. "I only charged \$10 for banging on the pipe. The other \$9,990 was for knowing where to hit it."

- 5 [00332] 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, or more generally, any skilled intelligent entity) 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
- 10 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.

- [00333] 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 present technology
- 15 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

- [00334] The second point (in Section 5.2) is that the expertise gathered via a human attention can be
- 20 re-used. In the plumber story, each time the expert plumber travelled to a new factory with a clogged pipe, he could re-sell 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 re-sell 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
- 25 be solved by an AI that learned from him. People are mortal and limited in their ability to learn and re-use 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
- 30 plumber's attention and knowledge that was acquired by the AI. Given this ability of AGI (in the present technology) to learn knowledge once and re-use repeatedly with essentially zero marginal cost, what then is the true value of human attention?

[00335] 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.

[00336] AI generally, and the present technology specifically, enable this novel approach to monetization.

10 5.2c Attention arbitrage opportunities

[00337] 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.

[00338] 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.

[00339] Similarly, today, few if any companies recognize that human attention is still under-valued by several orders of magnitude given advances enabled by new technology. As the advancements in AI, including the present technology, begin to show how attention can be monetized at much higher rates, the cost for ads – the type described in the present technology at least – will begin to rise to reflect the newly discovered ways to create value with human attention, as described above.

[00340] 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. Note that an arbitrage opportunity also exists, more generally, for any intelligent entity whose attention is being employed in low-value tasks (e.g. watching ads) rather than using the same amount of attention and intelligence in high-value tasks (e.g. solving valuable problems).

- 5 [00341] As discussed above, the value of creating smarter AI is hard to over-estimate. 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.
- 10 [00342] 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 present technology, 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 present technology, combined with existing Ad targeting and purchasing technology, enables extremely profitable Ad
- 15 arbitrage opportunities that were previously impossible to achieve.

5.3 Opportunities to Enhance AI safety

- [00343] 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
- 20 opportunities afforded by the present technology are in the area of increasing AI safety.

- [00344] 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 the feedback to AI systems as part of the final stage of their training. Certainly, the present technology can use online ads to harness that attention and focus it
- 25 on RLHF activities, resulting in safer systems. But this is just one use of the present technology to improve safety via what has become a conventional approach to AI safety.

- [00345] An even more innovative approach is to use the present technology 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
- 30 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.

[00346] 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 is inherent in the very operation of the system.

[00347] The applicant's present technology 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 AAAs and personalized PSIs.

[00348] 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 the 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 AAAs, that will enable AGI to have a robust and human-aligned value system.

[00349] 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 present technology discloses how the existing technology and apparatus of online Ad systems can be co-opted and used to acquire the necessary attention and focus it in 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 in the present technology is using human attention to make AGI and SuperIntelligence 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 Representative and Valid Sample of Human Values

[00350] 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 that 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 do not only have to be in the form of 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 technology of FIGS. 1-13 and disclosed in the applicant's pending patents, can be used to engage humans in problem solving scenarios with ethical considerations as simple of 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.

[00351] Note that in the future, when intelligent non-human entities become autonomous and potentially determine their own values and ethical preferences, these same methods (e.g. polls and surveys) that garner ethical input from humans might also be used to garner input from autonomous intelligent entities. However, human values must take precedence if we desire the SuperIntelligent systems of the future to remain aligned with human values.

[00352] The present technology, focused on identifying and capturing human attention to solve valuable problems, can be used for ethical problems as well as technical problems. Indeed, in the preferred implementations of the AGI technology, 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

[00353] One of the advantages to the approach to ethics and safety reflected in the AGI technology 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, all times, and 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.

[00354] 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 that is widely accepted today

(treatment of animals comes to mind) may someday appear as barbaric, cruel and intolerable. Democratic ideals that are cherished in one country are deemed a threat in another, even though both countries exist at the same time and hold values that are cherished by their respective cultures.

[00355] 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 rule can hope to keep pace with both the sheer scope of situations and the rapid rate of change and emergence of new situations. The only sustainable approach has to involve mechanisms to recruiting, on-demand and as needed, humans to weigh in on the ethics of various situations. The present technology's means of recruiting targeted humans, on-demand, in real-time, is ideally suited to this task.

[00356] Eventually, when AGI becomes vastly more intelligent and much faster than humans at cognitive activity, the same mechanisms that are described in the present technology with application to human attention, can also be applied to the attention of any intelligent entity (whether human or non-human). The present technology 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 in a second.

[00357] 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 are considerations that must be accounted for in the very design of the system. Scalability is not a nice-to-have feature. It is essential for the sustainable human safety.

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

[00358] Online advertising is becoming an increasingly difficult business as competition grows more intense for the limited resource 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. In 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?

[00359] 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 steppingstone 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.

[00360] 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
5 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
10 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, the intelligent entities will increasingly be customized and personalized AI agents (e.g., AAAs) as described by the applicant in his and other previous pending patents. The systems and methods for harnessing and coordinating and increasing this
15 intelligence in an overall AGI system that is capable of SuperIntelligent performance, has 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 creation of AGI and higher levels of intelligence (per Jensen Huang's vision) generally?

[00361] The applicant sees the opportunities as including, without limitation:

- 20 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 cloud and datacenter infrastructure that previously existed to support search, content, and advertising platforms to support the training of, and implementation and use of, the new
25 intelligent entities that are 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 the exactly the right time to enable efficient and effective SuperIntelligent cognitive behavior, including problem solving.

30 [00362] 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 just 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

- 5 [00363] Some companies, such as Google, are in the enviable position of not only having sophisticated online Ad capabilities and reach, but also 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
- 10 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 that like Google which want to dominate the competition but in a way that is safe for humanity.

15 5.6 Benefits to Humanity from Wide Deployment of the Present Technology

- [00364] As the applicant, has described in this and other patents, AGI and particularly SuperIntelligent AGI, represents 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
- 20 it and therefore, he might as well just accept the possibility and enjoy the innovative time in which we live. While I respect Elon's positive attitude and willingness to accept that which we cannot control, Elon (and others who share this view) are making a grave error.

- [00365] 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
- 25 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.

- [00366] Worse than that, they do not 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
- 30 something, and that something is growing 10X more powerful every six months, one natural reaction is to be afraid of it and to 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.

[00367] The applicant sometimes describes the current state of AI as being 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 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!

[00368] Geniuses like Musk do not 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. But 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 path forward with a limited window in which to act to design safe AGI and SI.

[00369] What are the key facts related to the path forward? Well:

- 1) AI safety must be designed into the system, not tested-in as currently practiced.
- 2) The design cannot have safety or “alignment” localized to one part of the system where it can easily be re-programmed or changed. That is, a “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 that are required for advanced cognition. As with living organisms, which developed these perceptual systems first in primitive brains, so too, AI has developed these components of advanced intelligence first. But sequential cognition, e.g. problem solving, planning, 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 great 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 and 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 has not 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 common denominator and representation shared by both human and AI intelligent entities. This common architecture allows intelligent AI entities to learn from the 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 entity 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 satisfied using the inventive approaches, systems, and methods of the applicant.

[00370] Given that the current course of AI development is blind and Ad hoc compared to transparent and deliberate design of the present technology, and given that the present technology 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 present technology 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 a significant benefit of the present technology to humanity indeed! 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 System and Methods

[00371] The present technology 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. The present technology is fundamentally about leveraging existing online Ad technology in the service of increasing the intelligence of AI systems or whatever type, existing and in the future.

- 5 **[00372]** A primary insight underlying the present technology is that human attention is being under-utilized 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 which operate 24X7 at much faster speeds than human intelligence can then
- 10 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. **[00373]** A final point is that as non-human intelligent entities develop, harnessing attention and expertise from these intelligent entities will become increasingly valuable as well. Just as there are worse and better uses of human attention, there will also be less and more valuable uses of the
- 15 attention and expertise of intelligent entities generally, as well as arbitrage opportunities associated with this fact. Every element of the current inventive technology that relates to the superior capture and monetization of human attention also applies more generally to all intelligent entities. Currently humans watch the ads, but the in future AI agents will also. Currently people are familiar with the idea of paying humans for expertise and attention, but in the future the notion of doing the same for
- 20 non-human intelligent entities will also become commonplace. For illustration of the technology, in this application generally and especially in the sections that follow, we describe the technology mainly in terms of application to humans with the understanding that it should also be clear that the technology also applies to any intelligent entity.

25 6.1 Overview Main systems components

[00374] The present technology, 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)
- 30 • Methods for capturing human attention as opposed to requiring the human agents themselves
- Methods for accessing Human Agents
- Systems and methos for implementing a Human (or Intelligent Entity) Attention Spot Market

- Systems and methods for capturing human (or intelligent entity) attention via Online Ad Units
- Systems and methods for identifying and accessing human (or intelligent entity) agents via Online Ad Units
- 5 • Compensation and payment mechanisms
- Reputational Mechanisms
- System and methods for supporting Problem Solving within Ad Units
- System and methods for Problem Solving outside of AD Units
- Feedback mechanisms to improve online Ad targeting
- 10 • 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
- 15 • 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
- 20 • Methods to support the Integration of Realtime and asynchronous data feeds

[00375] 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

- 25 [00376] 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
- 30 network have been described in detail in previous PPA and PCT applications that are incorporated by reference into this application.

[00377] 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 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 using 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.

[00378] 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 present technology 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, increases.

[00379] Thus, the present technology 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.

[00380] The value of such a SuperIntelligent system is extremely high enabling the present technology 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 (AAAs)

[00381] As has been described in earlier cited PPAs and PCTs, an important aspect of the present technology 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.”

[00382] In the future, depending on the direction in which the laws surrounding AI develop, these customized AI agents, referred to as AAAs 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 AAAs 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 is his primary concern, to the degree that “human rights” exist, such rights should be extended to all sentient beings at least to varying degrees. But these are a questions for future debate. As of the writing of this disclosure, the most
5 advanced AIs and AAAs are still considered technology, or “tools” of human owners and their level of intelligence is still inferior to most humans in most areas.

[00383] With regard to the present technology, the important point is that AAAs need data and information in order to be customized and personalized. This data, at the moment, is mainly dependent on other intelligent entities (currently humans) providing it. Humans can provide data to
10 train AAAs by directly applying their intelligence to answering questions, solving problems, or instructing AAAs. 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 AAAs), as when, for example, they “surf” the internet or engage in online activity that leaves a behavioral data trace.

15 [00384] 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 AAAs in specific areas where gaps in the AAAs’ 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.

20 [00385] The present technology can help because online Ad technology can be used to target humans (or other intelligent 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
25 that makes all of 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.

30 [00386] It Is our unique experiences that make us unique. And so too with AAAs. These unique experiences, as the VISA commercial says, are “priceless.” Similarly, the unique knowledge, data, and information that customizes and AAAI and makes is different from another is the source of the majority of value of that AAAI. Therefore, obtaining the unique knowledge in an efficient and effective way is extremely valuable. Adapting and improving online Ad technology as described in

the present technology represents and novel and useful means to gather and use this most valuable differentiating knowledge and information.

[00387] While means for customizing AAAs 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 is 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
5 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
10 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
15 mechanism, can be facilitated by the online Ad unit technology.
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.
- 20 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
25 auditable (potentially block-chain based) record of solution steps which enhance 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
30 solving or expertise gathered via the online Ad unit technology 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 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

- [00388] A distinction can be made between human beings and their attention (or intelligent entities themselves and the attention of such entities). 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 activity of the user 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.

[00389] 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.

[00390] However, actually, the value resides not with the humans (or intelligent entities) themselves but with their attention, and specifically the unique intelligence that can be produced via human attention. Intelligence is the important thing – not humans (or entities).

[00391] 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.

[00392] 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 present technology, 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.

[00393] 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 make pleasant conversation. Anyone can lug a toolbox. It is only the relatively few seconds when the plumber looks at the specific drainpipe that is clogged and uses his/her/their specialized knowledge about pipes and tools, that the specific attention of the plumber is needed as opposed to generic attention from any human.

[00394] 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 his/her/their hour applying just the 60 seconds of specialized plumbing knowledge 60 times instead of once, imagine how much more he/she/they could bill! That same idea – namely of identifying and isolating the

value-added human attention and knowledge – underlies the tremendous value that the present technology can provide.

5 [00395] 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 do not. That specific knowledge is what the online Ad technology can target, obtain, and then use to customize "plumber AAAs" or AGI that possess the ability to solve plumbing problems.

10 [00396] 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 present technology. Online Ad technology that helps secure just the right information at just the right time (which can then be re-used in millions of AI systems) should command prices 10X, 100X, or even 1000X what a "regular" online Ad would
15 generate.

6.5 Human Agent / Experts database

[00397] 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
20 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. But 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?

25 [00398] 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 ever 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
30 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.

[00399] 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 the present technology, 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.

10 [00400] 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 present technology to build a more comprehensive database of human agents or experts who can be contacted, using the present technology 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:

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 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 be used to continually improve the AI system's accuracy and efficiency based on expert feedback and results.
- 10 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.
- 15 7. Indexing: Indexing improves the speed of data retrieval operations by efficiently locating data without having to search 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.
- 20 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.
- 25 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.
- 30 10. Data Visualization Tools: Data visualization tools help in representing data in graphical formats. These tools can be used to 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 serve as I/O to AI systems as well.
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 present technology.

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.
- 5 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 present technology.
- 10 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 speed and efficiency of parts of the present technology.
- 15 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 with the present technology for secure transactions.
- 20 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. Realtime aspects of the present technology are also discussed in Section 6.24, where these systems also have relevance for implementation.
- 25 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.
- 30 18. Data Compression: Data compression reduces the size of the database. It can be particularly useful for storing large volumes of expert-related data, such as video interviews or detailed profiles, efficiently. 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 (or intelligent entity) Attention Spot Market

[00401] 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 how to solve similar problems forever in the future, how does one determine a fair price for that attention?

[00402] The present technology suggests that market mechanisms are the most efficient and fair ways we currently know to arrive at fair prices, not only for stocks and commodities, but for human attention as well. The present technology envisions multiple intelligent entities bidding on human (and non-human) attention using an “attention spot market.” In such a market, the humans can set the going rate for their time, allocating it on a first-come first-serve basis or via other means that are well known in the art of market mechanisms. As noted above, that what is true of humans, is true of intelligent entities generally. The current technology enables not only humans, but any intelligent entity to monetize attention and expertise via mechanisms such as the Attention Spot Market. For purpose of illustration, and also because humans are the logical first application of the technology, we shall continue to use the terms like Human Attention Spot Market and describe the technology as it applies to humans, while continuing to emphasize our understanding that it can be applied to any intelligent entity as those entities develop in intelligence.

[00403] FIG. 15 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, or 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 credit-worthiness of entities issuing bonds and other securities, which ratings are important factors in determining the prices of such securities;

- b. Use of feedback mechanisms as detailed in Section 6.15.
 - 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
- 5 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
- 10 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
- 15 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 and matched with Ask Prices, and when Bid Price – Ask Price, a transaction occurs that is binding on the purchaser and the seller of attention.
- 20 **[00404]** Further, variations of implementations of the basic elements outlined above and in FIG. 15 are 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.
- 25 **6.6a Direct Exchange Platform Implementation**
- [00405]** 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.
 - 30 2. Demand/Supply Listing: Sellers list their available time and expertise areas, while buyers list their needs and the time slots they are 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.
- 5 6. Feedback System: After each session, buyers and sellers rate each other, influencing future pricing and matching.

[00406] FIG. 16 is a diagram showing the flow from registration to transaction, highlighting the dynamic pricing and matching engines.

10 6.6b Auction-Based Marketplace Implementation

[00407] An alternative implementation of the Attentional Spot Market could include, without limitation, the following process steps:

[00408] Process Steps:

1. User Registration: Similar to the direct exchange, both buyers and sellers create profiles detailing their needs or expertise.
- 15 2. Auction Creation: Sellers create auctions for their time, specifying the minimum bid or use 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.
- 20 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.
- 25 [00409] FIG. 17 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

- [00410] The two alternative implementations of the attention spot market described in 6.6a and 30 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 are 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.
- 5 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.
- 10 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 6.11) 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.
- 15 8. Time Slot Segmentation: Prices vary by time slot, with peak times priced higher. Allows for price optimization based on demand patterns. Especially for real-time access to attention, this approach may be useful. Just as with Uber, electricity demand, or any other limited commodity, there may be certain times when the attention is more needed and more valuable than others.
- 20 9. Subscription Access: Buyers pay for a subscription for access to 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.
- 25 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 be likely used as a form of compensation for the human experts.
- 30 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 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 are they to sell their time) or the clients (how urgently do 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.

[00411] Each of these methods offers a unique way to manage the dynamic exchange of human (or intelligent entity) 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 (or intelligent entity) Experts

[00412] Here is one exemplary implementation of the design, appearance, technical operation, interaction process for an online Ad unit designed to help build a database of human (or intelligent entity) experts.

[00413] a) Design and Appearance:

[00414] 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."

[00415] b) Technical Operation:

[00416] 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.

5 **[00417]** c) Interaction Process:

1. 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.
- 10 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.

15

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

[00418] Here is one exemplary implementation of the design, appearance, technical operation, interaction process for an online Ad unit designed to capture expertise directly from human experts.

[00419] a) Design and Appearance:

- 20 **[00420]** 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.

25 **[00421]** b) Technical Operation:

- [00422]** 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 rewards points, redeemable for various incentives. The system would also track contributions to identify top contributors for potential future engagement.
- 30

[00423] 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, FIGS. 2 & 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.

5 **[00424]** c) Interaction Process:

1. The expert notices the Ad presenting a challenge related to their field of 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.
- 10 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 offer opportunities for user to click out of the Ad and into the more featured
- 15 problem solving environment with multiple intelligent entities as described in Section 4.5 and in previous cited PPAs and PCTs.

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

[00425] Here is one exemplary implementation of the design, appearance, technical operation, interaction process for an online Ad unit designed to both capture expertise directly from human

20 experts and also build a database of human experts.

[00426] a) Design and Appearance:

[00427] 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

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

[00428] b) Technical Operation:

[00429] Upon providing a solution, users would be redirected to a form to enter their professional details and join the expert database. This process would not only capture the immediate contribution

30 but also 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.

[00430] c) Interaction Process:

1. An expert engages with the Ad by solving a presented problem.

2. After submitting their solution, they are invited to 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, where they are 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 offer opportunities for user to click out of the Ad and into the more featured problem solving environment with multiple intelligent entities as described in Section 4.5 and in previous cited PPAs and PCTs.

6.10 Exemplary Compensation Methods and Processes

- [00431]** Section 4.5, FIGS. 2 & 3, and previously cited PPAs and PCTs describe compensation and payment mechanisms in the context of a problem solving network that supports AGI. However, it is also desirable to compensate human experts within an Ad unit should they decide to work only within the Ad unit. In this case, without limitation, one or a combination of the following ten exemplary methods and processes may be used within the Ad unit, as well as potentially within the AGI system previously described. Compensation would be primarily for the humans or intelligent entities providing their attention, information, expertise, or knowledge.

1. Direct Monetary Compensation via Digital Wallets

[00432] Compensation Method:

- [00433]** 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.

[00434] Process:

1. After a user submits their expertise through the Ad unit, an algorithm assesses the quality of the contribution.
2. The user is notified of their compensation amount based on this assessment.
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 chosen digital wallet.

5. The user receives a confirmation of the payment via email or notification from the digital wallet service.
2. Cryptocurrency Rewards
- 5 [00435] Compensation Method:
- [00436] Contributors are rewarded in cryptocurrency, which allows for instant, global payments without the need for traditional banking infrastructure. This could involve popular cryptocurrencies like Bitcoin or Ethereum, or a proprietary token created specifically for the platform.
- [00437] Process:
- 10 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.
 - 15 5. Users receive a notification of the transaction completion.
3. Gift Cards and E-Vouchers
- [00438] Compensation Method:
- [00439] Users earn points for their contributions, which can be exchanged for gift cards or e-
- 20 vouchers for popular online retailers, such as Amazon or eBay.
- [00440] 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.
 - 25 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.
4. Professional Development Opportunities
- [00441] Compensation Method:
- 30 [00442] In lieu 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 field of expertise.
- [00443] Process:
1. Contributions are evaluated for their impact and relevance.

2. Based on their contributions, users are offered access to select professional development resources.
3. Users select their preferred resource and provide necessary details for enrollment.
4. The platform arranges access to the resource, either by enrolling the user directly or providing
5 access codes.
5. Users receive confirmation and instructions on how to access their chosen resource.

5. Exclusive Content Access

[00444] Compensation Method:

- 10 **[00445]** 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.

[00446] 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.
- 15 3. Users select the content they wish to access and confirm their choice.
4. The platform unlocks access to the content for the user, often through a unique access link or code.
5. Users are notified and provided instructions on how to access the content.

20 6. Recognition and Awards

[00447] Compensation Method:

[00448] 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.

[00449] Process:

- 25 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
30 related communications.
5. This recognition serves as a professional accolade, enhancing the contributor's reputation in their field.

7. Referral Bonuses

[00450] Compensation Method:

[00451] Users are incentivized to refer other experts to the platform, receiving bonuses for each successful referral that results in a contribution.

[00452] Process:

- 5 1. Users are provided with 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 in the form of direct payments, points towards rewards, or other incentives.
4. The platform tracks referrals and contributions to ensure accurate compensation.
- 10 5. Referrers are notified of their bonus and provided with details on how to claim it.

8. Subscription Credits

[00453] Compensation Method:

- [00454]** For platforms that offer subscription services, contributors can receive credits towards their subscription fees, effectively reducing or waiving their costs. Credits can also be for use of the attentional spot market, the online Ad technology, or the AGI system generally.
- 15

[00455] 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.
- 20 3. Users are notified of the credits earned and the resulting 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.

9. Physical Merchandise

25 **[00456]** Compensation Method:

[00457] Users can choose from a range of branded merchandise, or relevant products related to their field of expertise as a form of reward.

[00458] Process:

1. Contributions are assessed, and users earn points based on their input.
- 30 2. Users can browse a catalog of available merchandise and select items based on the points they have accumulated.
3. Upon selection, users provide shipping details.
4. The platform processes the order and dispatches the merchandise to the provided address.
5. Users receive the merchandise as a tangible reward for their contributions.

10. Conference and Event Sponsorships

[00459] Compensation Method:

5 **[00460]** Highly active or valuable contributors can receive sponsorships for professional conferences or events, covering registration fees, travel, or accommodation.

[00461] Process:

1. Contributions are evaluated for their impact, with particular emphasis on contributors demonstrating consistent, high-quality involvement.
2. Eligible contributors are offered sponsorships for upcoming industry events relevant to their
10 expertise.
3. Interested contributors accept the sponsorship and provide necessary details for registration and travel arrangements.
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
15 minimal personal expense.

6.11 Reputational Mechanism

[00462] Reputation metrics should be calculated for all human experts and intelligent entities. Reputational metrics are important elements in the methods for deciding which problems or other
20 cognitive work to assign or request from which entities, and also in 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 more sophisticated matching to the requirements of clients or the task. For
25 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.

[00463] First, we list, without limitations, some of the potential dimensions / metrics of interest with
30 regard to 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 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 problems.
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 in identifying experts for mentorship roles or collaborative projects. Implementation can be via the present technology where the task of rating a peer is task 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 linked in profiles, and other public information. Influences assignment to projects requiring recognized expertise.
8. Innovation Score: Assesses creativity and the ability to generate novel solutions. Evaluated through peer reviews and client feedback (See 6) on the novelty. 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: Measures 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 based on successful application of 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 measures of agility.

14. Conflict Resolution: Assessed by observing outcomes of conflicts the expert was involved in. 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.
- 5 15. Project Management Skills: Evaluated based on the ability to plan, execute, and close projects effectively. Important for roles with project oversight responsibilities. Metric might be the number of project management tasks assigned to expert and 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
10 response time and metrics associated with items 1-5 above.
17. Efficiency: Assesses the amount of 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
15 longer term project performed outside of the Ad unit (See Section 6.13).
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 longer term project performed outside of the Ad unit (See Section 6.13).
20. Client Retention Rate: The percentage of clients who return or continue working with the expert.
20 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.
- 25 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 longer term project performed outside of the Ad unit (See Section 6.13).
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
30 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 useful solutions, options generated by humans can be rated on these dimensions by peers, with usefulness also correlated to problem success.

[00464] Reputational metrics, to the degree possible, should be automatically recorded and updated. For example, time-to-solution and steps-to-solution for problem solving is easy to automate, as is the percent 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 already exist and are in use, can also be used as part of the evaluation of a public reputation.

[00465] All of these dimensions of reputation can be used and updated in the present technology 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.

6.12 Problem Solving within Ad Unit

[00466] The present technology 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 so as to improve the capabilities of the overall AGI network. This AGI/ SI system has been described in detail in previously cited PPAs and PCTs, in Section 4.5, and in FIGS. 1-13. The general problem solving process is described in FIGS. 2, 4, & 10.

[00467] 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 present technology where one use of the online Ad unit is to solicit expertise and knowledge that help solve problems and also be used to train AI, one exemplary method to implement problem solving within an Ad unit is shown in FIG. 18. 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.
 - 10 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.
 - 15 e. The amount of credit or payment that the user of the online has accumulated based on his/her/their input.
 - f. 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.
- 20 6) If the problem / sub-problem is solved, the user is credited for his/her/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 stops 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.
- 25 8) The solution learning process of the AGI system that is 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 is acting as a valuable means of not only solving a particular problem, but also as a valuable means to increase the knowledge, intelligence and capabilities of the overall AGI system.
- 30

[00468] 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.

[00469] 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 FIGS. 2, 4, & 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 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.

5 [00470] 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 present technology. 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
10 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 cannot 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
15 superior to average unless exceptionally rich (in expertise) data is used to train.

[00471] The present technology of using online Ad units to identify and acquire information from humans that possess such superior expertise is therefore particularly helpful (and valuable) to address 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
20 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.

[00472] A second clarification is important for the future in which AI agents routinely act on behalf
25 of human users to represent their interests online. In such scenarios, which are becoming increasingly common, and which will likely represent the majority of internet traffic in the future, it is not humans who will predominately be exposed to online ads, but rather AI agents.

[00473] 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
30 meant in the most general, powerful, and preferred implementations of the present technology 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.

[00474] 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 the present technology, and other means, is ethical information since that depends, to put it poetically “on the human heart,” and not on intelligence alone.

10 6.13 Problem Solving Outside of Ad Unit

[00475] 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 a YES / NO response to an ethical question, or perhaps 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.

[00476] Therefore, the present technology includes all of the problem solving methods and systems described in Section 4.5, in the associated FIGS. 1-13, and that has been detailed in the cited PPAs and PCTs that are focused on such systems and methods.

[00477] 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 user’s opinion 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 present technology, targeting the online Ad unit to experts that 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.

6.14 Feedback Mechanism to Online Ad Targeting

[00478] 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 was revealed to the online Ad vendors, then prices would rise accordingly.

[00479] However, in the preferred implementation of the present technology, 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 directed 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.

[00480] FIG. 19 illustrates the basic feedback mechanism for Ad targeting, 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 the present technology) 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 6.11, 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 with regard to

which information, knowledge, or expertise, was most effective at meeting or maximizing problem solving and/or AI training objectives.

- 5 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.

6.14a Keys to AI Dominance

10 [00481] Google/Alphabet, Meta, Amazon, Apple, and many other large technology companies that sell online advertising, also are extremely focused on gaining competitive advantage in the field of AI agents and AGI. A hitherto unrecognized advantage that these companies and organizations possess, if they use the present technology, 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
15 targeting capabilities in combination with the present technology 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).

20 [00482] Given that the means to obtain high-quality training data is rapidly becoming the bottleneck in the AI “arms race,” the present technology 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
25 the objective of delivering the maximally useful information, from an AI agent training perspective, for increasing the intelligence of their AI systems.

30 [00483] It is difficult to overstate the magnitude of this advantage. Consider that 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.

[00484] 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, and it requires 1% of the computing power (chips) to train the AI agents as well!

- 5 [00485] 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 present technology, to accelerate the learning and intelligence of their AIs, without increasing their need for chips or computational power. Of course, if they can do both – obtain
10 higher quality data AND lots of computing power – that is best. But as described, better data with less compute can achieve superior results to mediocre data with more compute. In short, certain organizations are sitting on the keys to AI dominance without realizing it, and the present technology, together with those detailed in previously cited PPAs and PCTs, are the keys to leapfrogging the competition!

15

6.15 Feedback Mechanism / Process for Attention Spot Market

- [00486] 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 its quality and usefulness. As described above, objective, automated, and/or third-party methods may be used to try to group or
20 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 than there is, for example, in grades of crude oil, corn, or soybeans. Fortunately, AI excels in recognizing patterns and making judgements such as credit-worthiness or selecting humans based on resumes. Further, the problem solving and AI training tasks have objective metrics of quality and success
25 (detailed above) that can be used to more precisely estimated the value of a given human's attention for a particular task than can be done, for example, with a bushel of soybeans.

- [00487] An important element, therefore, in accurately pricing human attention is knowledge both of the performance track record of the particular human(s) and the nature of the task that the human is being asked to perform. Using methods from KIT (described in earlier PPAs and PCTs) it is
30 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.

[00488] 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 tasks based on 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.

[00489] This feedback might be considered most helpful in objectively characterizing -- including without limitation type or category of human attention, knowledge, or information -- the nature of human attention that is offered for sale. However, ultimately, the appropriate price for human attention for a particular human is determined by the market price.

[00490] If a client overpays, and is dissatisfied, that dissatisfaction will be reflected in reputational metrics that are fed back to the market and 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 going forward.

[00491] One novel and useful feature of the feedback mechanism associated with the present technology generally, and the human attention spot market in particular, is the ability to record in great detail each and 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.

[00492] Thus, based on the feedback mechanism, the exact 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 the 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 and can help both parties make more profitable transactions on the spot marketplace compared to more commoditized markets. In short, humans are unique, and the present technology helps humans get paid fairly for their uniqueness while also helping clients avoid over-paying for commoditized attention that is of less value.

[00493] One exemplary implementation of the feedback process for the attention spot market includes the following steps, as also illustrated in FIG. 20:

1. Purchase of human attention is contracted on the spot market; price and other details of the contract are recorded.
- 5 2. Work is performed by human.
3. Metrics, which may include, without limitation, the reputational metrics of Section 6.11 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
5. All metrics are stored in (optionally blockchain-technology enabled) transparent and auditable records associated with the transaction.
- 10 6. 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
- 15 (including from particular individuals or categories of individuals) as related to particular types of tasks are updated.
7. 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
- 20 price at which transactions occur on the spot market.

6.16 Continuous Improvement Mechanisms

- [00494] Similar to the feedback mechanism and process for the spot market, described in Section 6.15, the effectiveness of all steps in the overall online Ad unit technology, 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 is 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.

30 6.17 Dynamic Arbitrage Process

[00495] 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 or 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 technology, 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 arbitrating 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.
- [00496] 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.
- [00497] 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. Use of the online Ad unit technology is therefore a catalyst for the growth of intelligence of these AIs.
- [00498] 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. Terrorist (humans or AIs) might pay extremely prices for specific expertise related to weapons of mass destruction, for example, and this must be prohibited.

6.18 Human Worker & Client Interfaces

- [00499] User interfaces for human workers and clients can include LLMs and natural language text or audio-based interfaces in which human 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.
- 5 However other, more specialized interfaces can be built as well.
- [00500] 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.,
- 10 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.

6.19 Automated AI / AGI Interfaces

- 15 [00501] In addition to using interfaces for humans (e.g. human clients or workers), the present technology 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.
- [00502] 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
- 20 and generation beyond the range of humans.
- [00503] 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.
- 25 [00504] 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.

6.20 Recursive Use of Problem Solving to Optimize Ad Targeting and System Efficiency

- 30 [00505] One of novel and extremely powerful aspects of the present technology is the ability to solve cognitive problems of any type. Thus, although Sections 6.14 – 6.16 describe specific approaches to use feedback and other specific methods for improving aspects of the present technology, a general way to optimize aspects of the present technology, 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 present technology to improve the present technology is a unique capability that exists with almost no other inventions, and which reflects the universality of the problem solving capabilities that underly the present technology, and that are described in this disclosure and previously cited PPAs and PCTs.

6.21 Safety and Ethics Checks

[00506] 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. Terrorist (humans or AIs) might pay extremely prices for specific expertise related to weapons of mass destruction, for example, and this must be prohibited.

[00507] 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, AAAs 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 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.

6.22 Regulation Compliance

[00508] The same processes described in Section 6.21 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 present technology 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 regulation compliance.

6.23 Universality of System and Methods Across Platforms and Cultures

[00509] It should be clear from the above discussion that the online Ad unit present technology 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 present technology.

[00510] One of the unique and innovative features of the present technology is that it provides a means to access human attention, knowledge, and expertise, from almost any humans, 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 present technology 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.

10 6.24 Collaborative and Cross Ad Unit Dynamic Coordination

[00511] Another novel aspect of the present technology 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.

[00512] For example, multiple online Ad units could simultaneously solicit the “next step” in a problem from multiple humans. The 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 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 present technology. 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 technology. 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.

[00513] 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.

30 6.25 Integration of Realtime and Asynchronous Capabilities / Data feeds

[00514] Another novel feature of the present technology 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.

[00515] 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.

[00516] Alternatively, some task might be done asynchronously outside of the Ad unit (e.g. via an email system), while other tasks might be done in real-time within the Ad unit. Generally, the present technology 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

[00517] The present technology 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

[00518] Consider the Use Case of a company engaged in the sale of online ads whose business model is to maximize the profits from displaying client, or whose expertise is serving as broker between the client and the ultimate online Ad provider. Companies or organizations which 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 that sell subscriptions and products but also derive some significant revenue from advertising, such as Amazon and Apple, are also relevant to this Use Case.

[00519] In the example that follows, the company that possess the systems and methods of the present technology is called “Company” and the clients that wish to purchased human attention in order to extract information, knowledge, and expertise are called “Clients.”

5 [00520] The following steps illustrate one preferred implementation that results in increased monetization of online advertising revenue for Company and delivers to Client valuable information, knowledge and expertise based on human attention.

- 1) Company builds databases containing information about users that is helpful for targeting specific types of online ads to specific users. The means for building the databases include, without limitation, Company’s existing methods and technology, purchasing user data and
10 information from others, and using the present technology’s methods (Section 6.7) for building a database of human experts, which can include their preferences and other information useful 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 technology can be used specifically
15 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 Company. Client communicates the desired categories of human attention or work, which may include requirements with respect to certain populations of humans that can be accessed due to the
20 universal nature of the present technology (Section 6.22).
- 3) At the time that the requirements are communicated, checks are done to ensure that the request does not violate regulations or ethical requirements (Section 6.21). The requirements can be communicated via existing means or via interaction with a human attention spot market (Section 6.6).
 - 25 a. If an attention spot market is used, then 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 Client’s needs can be improved via a feedback loop and processes (Section 6.15).
- 4) Company can work with the client to design creative content for interactive online ads that
30 capture attention and work from human experts (Section 6.8), and 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 6.12) or by linking to webpages or other interfaces that are optimized for this capture outside of the Ad unit (Section 6.13).

- a. Depending on the nature of the work, the interactive ads may require collaboration across multiple humans (Section 6.24) and/or integration of real-time or asynchronous data capabilities (Section 6.25).
 - 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 6.21 and FIG. 8).
- 6) Based on reputational metrics and methods (Section 6.11) 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 6.14, 6.15, 6.16).

10

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

[00521] 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.

[00522] These Companies are especially well positioned to succeed in the race to develop the most advanced forms of AI by leveraging their own 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 present technology 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 own online Ad capabilities and can invest in the expenses of displaying the present technology's online Ad units in order to rapidly and efficiently increase the intelligence of their own AI agents and systems, including, without limitation AGI and SI systems (See Section 4.5).

[00523] 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 6.14 - 6.16). 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 own 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 6.15).

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

- 5 **[00524]** 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 present technology 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).
- 10 **[00525]** An important variation of the basic implementation outlined in Section 7.1, is therefore to focus the online units on the problems of reviewing, ranking, voting, and otherwise providing human input on potential ethical choices which may have safety implications. These problems can include efforts to gather representative and statistically valid sample of human values and opinions (Section 5.3a), across many diverse cultures and geographies, leveraging the present technology's
- 15 ability to access humans across many platforms (Section 6.23). They can also include real-time oversight on ethical and safety issues (Section 5.3b) – an area where the ability to access many human opinions in parallel on the same issue and coordinate the responses (Section 6.24) are 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 6.25) is also important.
- 20 **[00526]** 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 present technology 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

- 25 **[00527]** The description of the present technology 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 attention of ANY intelligent entity, whether human or AI.
- 30 **[00528]** In the general case, the Human Attention Spot Market (Section 6.6) becomes an Intelligent Entity Spot Market. Since AI agent, AGI and SI systems will, in the future, act as agents for humans (and themselves) representing the interests of other (e.g., human) entity's 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 6.19) rather than the standard human interfaces (Section 6.18) to the various inventions.

[00529] Just as profiles for Ad targeting can be developed for individual humans, and then used to help target the requisite human expertise, so too, profiles for Ad targeting can be developed for individual AI agents, and then used to target the requisite AI expertise.

[00530] The same distinction between a human and the human's attention (Section 6.4) applies to AIs. An AI is different than attention from that AI. For the purposes of the present technology, 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.

[00531] However, just as a database for locating the humans possessing the relevant expertise can be built using methods of the present technology (Sections 6.7, 6.8, 6.9) and reputations can be compiled and associated with the humans (Section 6.11), and feedback mechanisms can be implemented to improve targeting of the humans and the present technology's operations (Sections 6.14, 6.15, 6.16) 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.

[00532] 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 present technology is concerned with efficiently locating and tapping the intelligence (including associated information, knowledge, and expertise) to solve problems and especially to train and improve the intelligence of other AIs.

[00533] 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 present technology, including the online Ad units and attention spot market, are 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 present technology in the future.

[00534] The problem solutions, and increases in AI intelligence, which are produced by the present technology's novel method of monetizing (human or non-human) attention is many times the value produced today by online ads, which is why the present technology is not only novel, but extremely useful and valuable.

[00535] FIG. 21 is a diagrammatic representation of a computer system **100** that is utilizable or implementable with the user's device and/or any peripheral component of the present technology. The applicant notes that numerals **102 – 104** in FIG. 21 can refer to elements found in both classical computing architectures and quantum computing architectures, and that the present technology can
5 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 problem solving efficiency of AI agents or intelligent entities utilizing quantum computing architectures.

[00536] The computer system **100** can be part of an example machine, which is an example of one
10 or more of the computers referred to herein and, within which a set of instructions for causing the machine to perform any one of or more of the methodologies discussed herein may be executed. In various example embodiments, the machine operates as a standalone device or may be connected (e.g., networked) to other machines. In a networked deployment, the machine may operate in the capacity of a server or a client machine in a server-client network environment, or as a peer machine
15 in a peer-to-peer (or distributed) network environment. The machine may be a personal computer (PC), a tablet PC, a set-top box (STB), a personal digital assistant (PDA), a cellular telephone, a web appliance, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine.

[00537] The computerized system **100** can include one or more processors **102**, storage devices **106**,
20 and communication devices, as well as software components or instructions **104** for providing a platform for users to interact with and train/tune the LLMs. The computing capabilities may be stand alone or may be cloud based. They may include cloud based AI development platforms that seamlessly offer “AI as a service” and they may include both hardware and software components.

[00538] The system also supports the ability for users to provide new data, or data that is unique to
25 them, for the LLMs to learn from. The processors **102** may be one or more CPUs, GPUs, chips specialized for ML, microprocessors, application processors, embedded processors, field-programmable gate arrays (FPGAs), or other hardware components capable of executing computer programs. The processors may be in communication with one another and/or with other components of the system. Further, any one of or any combination of the components of the system **100** can
30 communicate with each other via a bus **134**.

[00539] The storage devices **106** may include one or more hard drives, solid-state drives, optical storage devices, or other storage components. The storage devices may store the data that is used to train/tune the LLMs, as well as other data associated with the system, such as user accounts, system settings, and other data.

[00540] The communication devices may include one or more cellular modems **108**, Wi-Fi cards **110**, Bluetooth modules **112**, Network Interface Device **114**, or other components that enable the system to communicate with other systems, such as user devices, over a network or the internet.

5 [00541] The communication devices may also enable the system to communicate with other systems over a wireless or wired connection **116**.

[00542] The software components may include computer programs for providing a platform for users to interact with and train/tune the LLMs. The software components may also include computer programs for collecting, storing, and processing data that is used to train and/or tune the LLMs. The software components may also include computer programs for providing a user interface for users to
10 interact with the system.

[00543] The user interface **118** may include, without limitation, natural language interfaces, textual interfaces, and chatbot type of interfaces, a web-based user interface, a mobile application, an augmented reality application, a metaverse application, or other applications that allow users to interact with the system. The user interface may include features for allowing users to select the data
15 that they want to use to train/tune the LLMs, as well as features for allowing users to interact with and monitor the progress of the LLMs.

[00544] The system may also include one or more databases or data source, including without limitation vector databases, centralized databases, and distributed databases, for storing the data that is used to train/tune the LLMs, as well as other data associated with the system, such as user
20 accounts, system settings, and other data. The databases may be hosted on the system itself or on another system, including cloud based systems.

[00545] The system may also include one or more authentication systems for verifying the identity of users who use the system, as well as for providing secure access to the system. The authentication systems may include biometric authentication systems **122**, such as facial recognition or fingerprint
25 recognition systems, as well as other authentication systems, such as password-based authentication systems.

[00546] The system may also include one or more security systems for protecting the system from unauthorized access and for protecting the data that is stored on the system. The security systems may include firewalls, encryption systems, access control systems, single and multi-factor
30 authentication systems, and other security systems.

[00547] The system may also include one or more analytics systems for collecting and analyzing data associated with the system and/or the LLMs. The analytics systems may include machine learning algorithms and other algorithms for analyzing the data associated with the system and/or the LLMs.

[00548] Data visualization methods, including use of problem trees and other representations and data structures; use of statistical outputs, tables, graphs, text, speech, video, image and graphical outputs may be used for one way or di-directional communication between users and the system, and between multiple (human or AI) agents or LLMs using the system to interact with each other in
5 large or small groups.

[00549] The system may also include one or more monitoring systems for monitoring the performance of the system and/or the LLMs. The monitoring systems may include systems for monitoring the performance of the system, such as system uptime, and systems for monitoring the performance of the LLMs, such as accuracy, speed, ethical compliance, reputation metrics, quality
10 metrics, and other metrics as discussed above or as are known in the art.

[00550] The system may include one of more of the architectures described above that enable one or more human or AI Agents or LLMs to engage in a variety of intellectual tasks including, without limitation, simple and complex and multi-step problem solving behavior with the system having all of the functionality and features previously described.

[00551] The system may also include one or more feedback systems for allowing users to provide feedback on the system and/or the LLMs. The feedback systems may include systems for allowing users to submit feedback on the system, such as bug reports, and systems for allowing users to submit feedback on the LLMs, such as suggestions for improving the accuracy or speed of the model.
15

[00552] The system may also include one or more management systems for managing the system and/or the LLMs. The management systems may include systems for managing the system, such as systems for managing the users and user accounts, and systems for managing the LLMs, such as systems for managing the data used to train and/or tune the model.
20

[00553] The system may also include one or more payment systems allowing users to pay for the use of the system and/or the LLMs. The payment systems may include systems for processing payments, such as credit card processing systems, and systems for managing payments, such as subscription management systems.
25

[00554] The system may also include one or more other components, such as support systems, reporting systems, and other components that are necessary for providing a platform for users to interact with and train/tune the LLMs.
30

[00555] The computerized system of the present technology enables users to interact with and train/tune LLMs based on data that is unique to the users. The components of the system described herein provide the necessary hardware and software components for enabling users to do so.

[00556] Further, while only a single machine is illustrated, the term “machine” shall also be taken to include any collection of machines that individually or jointly execute a set (or multiple sets) of instructions to perform any one of or more of the methodologies discussed herein.

5 [00557] The computer system 100 may further include or be in operable communication with a video display 120 (e.g., a liquid crystal display (LCD), touch sensitive display), input and/or output device(s) 130 (e.g., a keyboard, keypad, touchpad, touch display, buttons, sonic, sensorial, etc.), a cursor control device 132 (e.g., a mouse), a drive unit 124 (also referred to as disk drive unit), and a signal generation device 128 (e.g., a speaker). The drive unit 124 can include a computer or machine-readable medium 126 on which is stored one or more sets of instructions and data
10 structures (e.g., instructions 104) embodying or utilizing any one of or more of the methodologies or functions described herein. The instructions 104 may also reside, completely or at least partially, within the memory 106 and/or within the processors 102 during execution thereof by the computer system 100. The memory 106 and/or the processors 102 may also constitute machine-readable media.

15 [00558] Still further, the computer system 100 can be in operable association or communication with any type of multi-modal input and/or output 130 that address the human senses, as well as I/O technology that extends beyond the range of normal human perception. Such as the ability to process invisible to humans, for example but not limited to, X-rays and information outside of the typical bandwidths of human perception, but not outside of AI perception using tools. Additionally,
20 the I/O technology can include very fast perceptions that are too fast for humans to perceive but which an AI entity could perceive, and very slow or faint perceptions (e.g., tiny seismic shifts occurring over years) that humans cannot perceive but which AIs could. Since any intelligent entity can be part of the present technology system described by FIG. 21, then it can be appreciated that any type of I/O that humans, and also AIs with much broader perceptual capabilities than humans,
25 can be utilized with the system 100.

[00559] The instructions 104 may further be transmitted or received over a network via the network interface device 114 utilizing any one of a number of well-known transfer protocols (e.g., Hyper Text Transfer Protocol (HTTP)). While the machine-readable medium is shown in an example embodiment to be a single medium, the term “computer-readable medium” should be taken to
30 include a single medium or multiple media (e.g., a centralized or distributed database, vector databases, and/or associated caches and servers) that store the one or more sets of instructions. The term “computer-readable medium” shall also be taken to include any medium that is capable of storing, encoding, or carrying a set of instructions for execution by the machine and that causes the machine to perform any one of or more of the methodologies of the present application, or that is

capable of storing, encoding, or carrying data structures utilized by or associated with such a set of instructions. The term “computer-readable medium” shall accordingly be taken to include, but not be limited to, solid-state memories, optical and magnetic media, and carrier wave signals. Such media may also include, without limitation, hard disks, floppy disks, flash memory cards, digital
5 video disks, random access memory (RAM), read only memory (ROM), and the like. The example embodiments described herein may be implemented in an operating environment comprising software installed on a computer, in hardware, or in a combination of software and hardware.

[00560] An example machine system of the present technology including the computer system **100** in combinational and/or operational use with components of the present technology. In the
10 exemplary, any or all of above described components can include a processor **102**, memory **106**, a network interface device **114**, a display **120**, an input device(s) **130**, **132**, and/or drive unit **124**.

[00561] According to one aspect, the present technology can include a system or method of utilizing an interactive online advertisement to train or increase knowledge of an Artificial Intelligence (AI) agent or system. An interactive online advertisement is populated utilizing problem information
15 relating to a problem provided to an AI agent or system. The interactive online advertisement is configured or configurable to acquire human attention from human users with content associated with the problem information or the problem. Input information acquired by the interactive online advertisement from the human users is received by the AI agent or system. The AI agent or system utilizes the input information to update the interactive online advertisement or provides the input
20 information for use in a problem solving process on the problem.

[00562] According to another aspect, the present technology can include a system for utilizing online advertising technology for increasing an intelligence of an AI agent or system. The system can include:

a computer system comprising: a processor, a computer-readable storage medium, and program
25 instructions stored on the computer-readable storage medium being executable by the processor to cause the computer system to:
provide the online advertisement unit to one or more user computer systems utilizing a network;
receive one or more inputs from the user computer systems participating in the online
30 advertisement unit, the inputs being providing by a human user utilizing the user computer systems respective, and the inputs being associated with any one of or any combination of solving the problem, solving a sub-problem of the problem, and advancing progress on the problem;
communicate the inputs to the AI agent or system;

utilize the inputs by the AI agent or system in a universal problem solving method to generate a solution to the problem or the sub-problem; and

train the AI agent or system with results from the universal problem solving method that are based on successful or unsuccessful solution attempts to solve the problem or the sub-problem, thereby increasing an intelligence of the AI agent or system.

5

[00563] According to yet another aspect, the present technology can include a method for online advertising technology utilized with an AI agent or system for increasing an intelligence of the AI agent or system. The method can include the steps of:

receiving problem information of a problem provided by an AI agent or system;

10

populating an online advertisement unit including the problem information;

providing the online advertisement unit to one or more user computer systems utilizing a network;

receiving one or more inputs from the user computer systems participating in the online advertisement unit, the inputs being providing by a human user utilizing the user computer systems respective, and the inputs being associated with any one of or any combination of solving the problem, solving a sub-problem of the problem, and advancing progress on the problem;

15

communicating the inputs to the AI agent or system;

utilizing the inputs by the AI agent or system in a universal problem solving method to generate a solution to the problem or the sub-problem; and

20

training the AI agent or system with results from the universal problem solving method that are based on successful or unsuccessful solution attempts to solve the problem or the sub-problem, thereby increasing an intelligence of the AI agent or system.

[00564] Some embodiments of the present technology can include a step of receiving advertisement specifications from a client to be used in the populating the online advertisement unit in combination with the problem information.

25

[00565] In some embodiments, the advertisement specifications can include any one of or any combination of advertisement content, demographic information, location restrictions for the online advertisement unit, advertisement budget, and metrics for determining a successful solving of the problem or the sub-problem.

30

[00566] Some embodiments of the present technology can include a step of billing the client based on cost per thousand impressions or clickthrough rate metrics.

[00567] In some embodiments, the problem information can include any one of or any combination of a current problem state of the problem, a sub-problem of the problem, a goal of the problem, and a sub-goal of the goal.

5 [00568] Some embodiments of the present technology can include a step of communicating to the AI agent or system one or more additional inputs from intelligent entities in combination with the inputs from the participating user computer systems, the intelligent entities is any one of or any combination of an additional human user utilizing an additional computer system, an additional AI agent or system, an Artificial General Intelligent (AGI) agent or system, and a SuperIntelligent (SI) agent or system.

10 [00569] In some embodiments, the updating of the online advertisement unit can include any one of or any combination of whether the inputs was accepted or rejected, whether a goal or sub-goal of the problem has been achieved, if the input was selected by the AI agent or system to advance problem solving, a description of a new problem state after applying the selected inputs, an amount of credit or payment that the intelligent entities has accumulated based on inputs from the intelligent entities, 15 and an additional request for new user input by way of the online Ad unit or a link that directs the intelligent entities to an interface for participating in the problem solving process on the problem.

[00570] Some embodiments of the present technology can include a step of crediting the intelligent entities if the problem or the sub-problem is solved, based on an amount of contribution by the participating intelligent entities.

20 [00571] Some embodiments of the present technology can include a step of soliciting, if the problem or the sub-problem is not solved, additional input from the human user of the participating user computer systems respectively.

[00572] Some embodiments of the present technology can include a step of compensating the human user when the human user exits the online Ad advertisement unit and stops participating in 25 the problem solving process.

[00573] In some embodiments, the training of the AI agent or system can further include a training method selected from the group consisting of differential privacy, federated learning, homomorphic encryption, synthetic data generation, secure multi-party computation, data anonymization, transfer learning, active learning, self-supervised learning, domain adaptation, reinforcement learning, few-shot learning, explainable AI, privacy-preserving record linkage, data augmentation, generative adversarial networks, crowdsourcing for data labelling, model personalization layers, knowledge distillation, and ensemble learning. 30

[00574] Some embodiments of the present technology can include a step of creating a database of intelligent entity experts, wherein the intelligent entity experts are any one of or any combination of

human user utilizing a computer system, an additional AI agent or system, an Artificial General Intelligent (AGI) agent or system, and a SuperIntelligent (SI) agent or system.

[00575] In some embodiments, the online advertisement unit can be provided to one or more of the intelligent entity experts from the database.

- 5 [00576] In some embodiments, the creating of the database can utilize a technology selected from the group consisting of relational database management system, NoSQL databases, data warehousing, data mining, machine learning algorithms, graph databases, vector databases, indexing, full-text search, blockchain, data visualization tools, application programming interfaces, data cleaning and preprocessing, cloud storage, caching, transactional database systems, real-time database systems, content delivery network, data compression, data encryption, and replication.

[00577] Some embodiments of the present technology can include a step of bidding on attention using an attention spot market.

[00578] In some embodiments, the attention spot market can include any one of or any combination of:

- 15 a means for the human (or intelligent entity) users with attention to sell to access a marketplace and specify seller information for sale and an ask price for the information;
a means for buyers of human (or intelligent entity) attention to access the marketplace and specify buyer information the buyer is willing to buy and a bid price for buying the buyer information;
20 a market mechanism for queuing the bid price and the ask price, including categories of the buyer and seller information, wherein each category has a market in the marketplace;
the market mechanism is configured or configurable to make the market in each category by market makers; and
the market mechanism is configured or configurable to match bid prices and ask prices, and a
25 transaction occurs that is binding on the buyer and the seller of the information.

[00579] In some embodiments, the attention spot market can include the steps of:

- buyers and sellers of human (or intelligent entity) attention and expertise register on a platform, the buyers providing details about buyer interests or expertise, the sellers providing details about sellers interests or expertise;
30 listing by the sellers available time slots and expertise areas, and listing by the buyers needs and time slots the buyers is interested in;
utilizing, by the platform, an algorithm to dynamically price human (or intelligent entity) attention and expertise based on supply, demand, and user ratings;

matching one or more of the buyers and one or more of the sellers based on requirements, availability, and price;

enabling transactions where the buyers pay for the seller time slots, and wherein the platform takes a commission; and

- 5 providing feedback where after each session, the buyers and the sellers rate each other, influencing future pricing and matching.

[00580] In some embodiments, the platform can utilize any one or combination of mechanisms selected from the group consisting of Dutch auction, reverse auction, sealed bid auction, open ascending price auction, fixed price with time priority, dynamic pricing based on ratings, supply-
10 demand curve adjustment, time slot segmentation, subscription access, freemium model, group buying, tiered expertise levels, flash sales, loyalty points system, hybrid auction, geographic pricing, behavioral pricing, escrow system, social influence pricing, and tokenized transactions.

[00581] In some embodiments, the attention spot market can include the steps of:

- 15 buyers and sellers of human (or intelligent entity) attention and expertise register on a platform, the buyers providing details about buyer interests or expertise, the sellers providing details about sellers interests or expertise;
- creating an auction by the sellers for seller time slots;
- placing one or more buyer bids by a buyer on time slots and expertise the buyer requires;
- closing the auction at a predetermined time or when the seller accepts a buyer bid;
- 20 paying the seller by the buyer of a winning bid, and the seller provides the attention or expertise based on the time slot, wherein the platform mediates an exchange and secures payment; and providing feedback where after each auction, the buyers and the sellers rate each other, affecting future auctions and visibility on the platform.

[00582] In some embodiments, the platform can utilize any one or any combination of mechanisms selected from the group consisting of Dutch auction, reverse auction, sealed bid auction, open ascending price auction, fixed price with time priority, dynamic pricing based on ratings, supply-
25 demand curve adjustment, time slot segmentation, subscription access, freemium model, group buying, tiered expertise levels, flash sales, loyalty points system, hybrid auction, geographic pricing, behavioral pricing, escrow system, social influence pricing, and tokenized transactions.

30 [00583] In some embodiments, the online advertisement unit can be configured or configurable to include material relevant to an area of expertise associated with the problem information.

[00584] In some embodiments, the online advertisement unit can be configured or configurable to, upon interaction by the human (or intelligent entity) users, would expand or redirect the human (or intelligent entity) user to a secure form on a landing page, wherein the form is configured to collect

user information and to provide questions to the human (or intelligent entity) user regarding to qualifications of the human (or intelligent entity) user.

5 [00585] Some embodiments of the present technology can include a step of validating the user information and categorizing the human (or intelligent entity) users in a database based on the qualifications.

[00586] In some embodiments, the online advertisement unit can be configured or configurable to include an interface allowing the human (or intelligent entity) users to input a solution to the problem or the sub-problem.

10 [00587] In some embodiments, the online advertisement unit can be configured or configurable to include any one of or any combination of elements of gamification, a scoring system, and a leaderboard of other human (or intelligent entity) users.

[00588] In some embodiments, the online advertisement unit can be configured or configurable to include text input fields or other interactive tools enabling the human (or intelligent entity) users to provide a solution to the problem or sub-problem directly within a space of the online advertisement unit.

15 [00589] In some embodiments, the solution can be stored in a database, and an algorithm is utilized to assess a quality and relevance of the solution based on predefined criteria.

[00590] In some embodiments, the human (or intelligent entity) users can receive instant feedback or rewards points based on the assessed quality and relevance of the solution, redeemable for various incentives.

20 [00591] In some embodiments, contributions by the human (or intelligent entity) users can be tracked to identify top contributors from multiple additional human (or intelligent entity) users for potential future engagement.

[00592] In some embodiments, the online advertisement unit can be configured or configurable to, upon interaction by the human (or intelligent entity) users, would expand or redirect the human (or intelligent entity) user to a secure form on a landing page, wherein the form is configured to collect user information and to provide questions to the human (or intelligent entity) user regarding to qualifications of the human (or intelligent entity) user.

25 [00593] Some embodiments of the present technology can include a step of validating the user information and categorizing the human (or intelligent entity) users in the database based on the qualifications.

[00594] Some embodiments of the present technology can include steps of:

analyzing the inputs from the participating human (or intelligent entity) users to determine a quality or relevance of the inputs, respectively; and

compensating the participating human (or intelligent entity) users based on the quality or relevance of the inputs.

[00595] Some embodiments of the present technology can include a step of providing digital wallet information by each of the participating human (or intelligent entity) users into a secure form, wherein the secure form is within the online advertisement unit or is provided by a linked platform.

[00596] In some embodiments, the compensation can be provided in any one of or any combination of cryptocurrency, electronic gift cards, electronic vouchers, subscriptions, access to online resources or content, recognition on online platforms, merchandise related to an expertise of the human (or intelligent entity) users respectively, and sponsorships for professional conferences or events.

[00597] Some embodiments of the present technology can include a step of incentivizing the participating human (or intelligent entity) users to refer other human (or intelligent entity) users to an online platform, by offering additional compensation for each referral that contributes to the platform, wherein the online platform receives the problem and provides the online advertisement unit.

[00598] Some embodiments of the present technology can include a step of receiving reputational metrics on each of the human (or intelligent entity) users.

[00599] In some embodiments, the reputational metrics can be multi-dimensional including any one of or any combination of timeliness, budget compliance, quality of work, solution success rate, client satisfaction, peer-rating of competency, external reputation, innovation score, communication skills, adaptability, leadership quality, technical proficiency, learning agility, conflict resolution, project management skills, reliability, efficiency, cultural fit, work ethic, client retention rate, feedback responsiveness, networking ability, mentorship and training, availability and responsiveness, problem-solving speed, and creativity and innovation.

[00600] In some embodiments, the reputational metrics for any one of the human (or intelligent entity) users can be automatically updated within the online advertisement unit.

[00601] Some embodiments of the present technology can include a step of updating the online advertisement unit based on the results from the universal problem solving method conducted on the problem or the sub-problem.

[00602] In some embodiments, the online advertisement unit can be configured or configurable to ask the human users to provide a description of how the problem is to be framed, and then the description is then translated into a language of the universal problem solving method using a natural language to problem solving language translator.

- [00603] In some embodiments, the online advertisement unit can be configured or configurable to ask the human (or intelligent entity) users to provide a description of intermediate goals towards a final solution of the problem, and wherein the intermediate goals are communicated to the AI agent or system for processing by the universal problem solving method.
- 5 [00604] Some embodiments of the present technology can include a step of requesting, by the AI agent or system, the human (or intelligent entity) users to provide safety or ethical information associated with the problem, the sub-problem or goals associated with the problem.
- [00605] Some embodiments of the present technology can include a step of utilizing the safety or ethical information in the populating or in updating the online advertisement unit.
- 10 [00606] Some embodiments of the present technology can include a step of modifying the online advertisement unit to ask the human (or intelligent entity) users to vote, rank or rate one or more potential goals, sub-goals or actions that have been proposed by any one of or any combination of the human users, and intelligent entities.
- [00607] Some embodiments of the present technology can include a step of requiring the AI agent or system to be identified as an AI agent so any inputs from the AI agent or system is excluded where only human opinions are sought.
- 15 [00608] In some embodiments, the online advertisement can be is provided by an online advertisement service provider that uses user information about the human (or intelligent entity) users to target the human (or intelligent entity) users with expertise relevant to the problem.
- 20 [00609] Some embodiments of the present technology can include a step of measuring and recording response rates for the participating human (or intelligent entity) users who contribute information within the online advertisement unit or by way of links to websites or interfaces outside the online advertisement unit.
- [00610] Some embodiments of the present technology can include a step of recording metrics related to a quality of information or solutions provided by the targeted human (or intelligent entity) users.
- 25 [00611] Some embodiments of the present technology can include a step of calculating correlations or statistical relationships on the user information to determine any one of or any combination of which factors have a highest impact on the metrics, which of the metrics is prioritized by the online advertisement service provider, and which of the inputs was most effective at meeting or maximizing problem solving or training of the AI agent or system.
- 30 [00612] Some embodiments of the present technology can include a step of performing arbitrage that can include the steps of:

estimating, by the online advertisement service provider, a value of a particular type of human (or intelligent entity) attention and knowledge with regard to the problem or sub-problem; estimating, by the online advertisement service provider, a cost of obtaining an amount of the human (or intelligent entity) attention required to complete the problem or sub-problem from the human users or intelligent entities by way of the online advertisement unit; and
5 determining if the estimated value exceeds the estimated cost by a predetermined or dynamic variable representing a profit margin, then purchase the human (or intelligent entity) attention at cost approximating that of the estimated value and sell the human (or intelligent entity) attention to clients, or use the human (or intelligent entity) attention to perform tasks
10 that create value approximating that of the estimated cost.

[00613] Some embodiments of the present technology can include a step of performing arbitrage in priority order of a largest arbitrage opportunities first, until a minimal acceptable arbitrage opportunity is reached.

[00614] In some embodiments, the online advertisement unit can include one or more interfaces
15 selected from the group consisting of Large Language Model and Small Language Model interfaces, natural language text interfaces, audio-based interfaces, text boxes, multimodal interfaces, dropdown lists, templates, dynamically sizing input and output areas, visual input or output devices, virtual reality devices, and augmented reality devices.

[00615] In some embodiments, the online advertisement unit can be multiple online advertisement
20 units with a first online advertisement unit being provided to a first set of the human (or intelligent entity) users for receipt of the inputs, and a second online advertisement unit being provided to the first set or a second set of the human (or intelligent entity) users and asking to vote on the inputs.

[00616] In some embodiments, the inputs can be provided to the AI agent or system in real-time for real-time utilization by the universal problem solving method.

[00617] According to still another aspect, the present technology can include a method of increasing
25 monetization of online advertising revenue for an online advertisement service provider by delivering to a client information based on human (or intelligent entity) attention. The method can include the steps of:

building, by an online advertisement service provider, one or more databases containing
30 information about human (or intelligent entity) users, the information is related to targeting specific types of online advertisements to specific human (or intelligent entity) users; communicating to the online advertisement service provider by a client a request to purchase human (or intelligent entity) attention, the request including requirements and desired categories of the human (or intelligent entity) attention;

performing checks on the request by the online advertisement service provider at a time that the requirements are communicated, the checks are configured or configurable to ensure that the request does not violate regulations or ethical requirements;

creating content for an interactive online advertisement unit that is configured or configurable to capture human (or intelligent entity) attention and work from human (or intelligent entity) users;

deploying, by the online advertisement service provider, the interactive online advertisement unit to a computer system of the human (or intelligent entity) users by way of a network;

allowing the interactive online advertisement unit to collaborate with multiple of the human (or intelligent entity) users and integrate with real-time or asynchronous data capabilities;

performing ethics and safety checks during a problem solving process on a problem provided by the client, the ethics and safety checks are configured or configurable to ensure that unethical or unsafe expertise is not part of the problem; and

improving the interactive online advertisement unit based on a feedback loop including on any one of or any combination of reputational metrics, and metrics related to the problem solving process or knowledge captured.

[00618] In some embodiments, the client can be any one of or any combination of an AI agent or system, a system that trains or customizes AI agents, and a system that provides solutions to problems.

[00619] Some embodiments of the present technology can include a step of determining if the database does not include a predetermined number of the human (or intelligent entity) users, or if additional human (or intelligent entity) users are required, then configuring the interactive online advertisement unit to acquire additional human (or intelligent entity) users.

[00620] In some embodiments, the requirements can be communicated by way of interaction with a human (or intelligent entity) attention spot market.

[00621] In some embodiments, the attention spot market can include any one of or any combination of:

a means for the human (or intelligent entity) users with attention to sell to access a marketplace and specify seller information for sale and an ask price for the information;

a means for buyers of human (or intelligent entity) attention to access the marketplace and specify buyer information the buyer is willing to buy and a bid price for buying the buyer information;

a market mechanism for queuing the bid price and the ask price, including categories of the buyer and seller information, wherein each category has a market in the marketplace;

the market mechanism is configured or configurable to make the market in each category by market makers; and

the market mechanism is configured or configurable to match bid prices and ask prices, and a transaction occurs that is binding on the buyer and the seller of the information.

5 **[00622]** In some embodiments, the attention spot market can include the steps of:

buyers and sellers of human (or intelligent entity) attention and expertise register on a platform, the buyers providing details about buyer interests or expertise, the sellers providing details about sellers interests or expertise;

10 listing by the sellers available time slots and expertise areas, and listing by the buyers needs and time slots the buyers is interested in;

utilizing, by the platform, an algorithm to dynamically price human (or intelligent entity) attention and expertise based on supply, demand, and user ratings;

matching one or more of the buyers and one or more of the sellers based on requirements, availability, and price;

15 enabling transactions where the buyers pay for the seller time slots, and wherein the platform takes a commission; and

providing feedback where after each session, the buyers and the sellers rate each other, influencing future pricing and matching.

20 **[00623]** In some embodiments, the platform can utilize any one or any combination of mechanisms selected from the group consisting of Dutch auction, reverse auction, sealed bid auction, open ascending price auction, fixed price with time priority, dynamic pricing based on ratings, supply-demand curve adjustment, time slot segmentation, subscription access, freemium model, group buying, tiered expertise levels, flash sales, loyalty points system, hybrid auction, geographic pricing, behavioral pricing, escrow system, social influence pricing, and tokenized transactions.

25 **[00624]** In some embodiments, the attention spot market can include the steps of:

buyers and sellers of human (or intelligent entity) attention and expertise register on a platform, the buyers providing details about buyer interests or expertise, the sellers providing details about sellers interests or expertise;

creating an auction by the sellers for seller time slots;

30 placing one or more buyer bids by a buyer on time slots and expertise the buyer requires;

closing the auction at a predetermined time or when the seller accepts a buyer bid;

paying the seller by the buyer of a winning bid, and the seller provides the attention or expertise based on the time slot, wherein the platform mediates an exchange and secures payment; and

providing feedback where after each auction, the buyers and the sellers rate each other, affecting future auctions and visibility on the platform.

5 [00625] In some embodiments, the platform can utilize any one of or any combination of mechanisms selected from the group consisting of Dutch auction, reverse auction, sealed bid auction, open ascending price auction, fixed price with time priority, dynamic pricing based on ratings, supply-demand curve adjustment, time slot segmentation, subscription access, freemium model, group buying, tiered expertise levels, flash sales, loyalty points system, hybrid auction, geographic pricing, behavioral pricing, escrow system, social influence pricing, and tokenized transactions.

10 [00626] Some embodiments of the present technology can include a step of providing a feedback mechanism associated the human (or intelligent entity) attention spot market, the feedback mechanism is configured or configurable to record each and every step in the problem solving process, and to create a vector track record of a performance of all the human (or intelligent entity) users on every task and sub-task.

15 [00627] In some embodiments, the vector track record can be implemented by way of blockchain technology to allow for precise reputations that are analyzed by an AI agent or system and converted into estimates of a value for each of the human (or intelligent entity) users for any of the tasks.

[00628] In some embodiments, the interactive online advertisement unit can be configured or configurable to capture the human (or intelligent entity) attention and work from within the
20 interactive online advertisement unit or by providing link to one or more webpages or interfaces that are optimized for the capture outside of the interactive online advertisement unit.

[00629] In some embodiments, the building of the databases can include the step of purchasing, by the online advertisement service provider, user data and information from remote sources.

[00630] According to yet still another aspect, the present technology can include a method for
25 online advertising technology utilized with an AI agent or system. The method can include the steps of:

populating, by an AI agent or system, an online advertisement unit including problem information provided to the AI agent or system;

30 providing the online advertisement unit to one or more user computer systems utilizing a network;

receiving one or more inputs from the user computer systems participating in the online advertisement unit, the inputs being providing by a human (or intelligent entity) user utilizing the user computer systems respective, and the inputs being associated with any one

of or any combination of solving the problem, solving a sub-problem of the problem, and advancing progress on the problem;

communicating the inputs to the AI agent or system;

utilizing the inputs by the AI agent or system or an additional AI agent or system in a universal

5 problem solving method to generate a solution to the problem or the sub-problem; and

utilizing a feedback mechanism that provides metrics to the AI agent or system to train the AI agent or system, the metrics assign credit or blame to the participating human (or intelligent entity) users and targeting mechanisms that directs the online advertisement unit to particular (or intelligent entity) human users.

10 **[00631]** Some embodiments of the present technology can include a step of measuring and recording response rates for the participating human users who contribute information within the online advertisement unit or by way of links to websites or interfaces outside the online advertisement unit.

[00632] Some embodiments of the present technology can include a step of recording the metrics
15 related to a quality of information or solutions provided by the targeted human (or intelligent entity) users.

[00633] Some embodiments of the present technology can include a step of calculating correlations or statistical relationships on the user information to determine any one of or any combination of which factors have a highest impact on the metrics, which of the metrics is prioritized by the online
20 advertisement service provider, and which of the inputs was most effective at meeting or maximizing problem solving or training of the AI agent or system.

[00634] Some embodiments of the present technology can include a step of performing arbitrage that can include the steps of:

estimating, by the online advertisement service provider, a value of a particular type of human
25 (or intelligent entity) attention and knowledge with regard to the problem or sub-problem;

estimating, by the online advertisement service provider, a cost of obtaining an amount of the human (or intelligent entity) attention required to complete the problem or sub-problem from the human users or intelligent entities by way of the online advertisement unit; and

determining if the estimated value exceeds the estimated cost by a predetermined or dynamic
30 variable representing a profit margin, then purchase the human (or intelligent entity) attention at cost approximating that of the estimated value and sell the human (or intelligent entity) attention to clients, or use the human (or intelligent entity) attention to perform tasks that create value approximating that of the estimated cost.

[00635] Some embodiments of the present technology can include a step of performing arbitrage in priority order of a largest arbitrage opportunities first, until a minimal acceptable arbitrage opportunity is reached.

5 [00636] While embodiments of the online advertising technology for AGI and SI have been described in detail, it should be apparent that modifications and variations thereto are possible, all of which fall within the true spirit and scope of the present technology. With respect to the above description then, it is to be realized that the optimum dimensional relationships for the parts of the present technology, to include variations in size, materials, shape, form, function and manner of operation, assembly and use, are deemed readily apparent and obvious to one skilled in the art, and
10 all equivalent relationships to those illustrated in the drawings and described in the specification are intended to be encompassed by the present technology.

[00637] Therefore, the foregoing is considered as illustrative only of the principles of the present technology. Further, since numerous modifications and changes will readily occur to those skilled in the art, it is not desired to limit the present technology to the exact construction and operation shown
15 and described, and accordingly, all suitable modifications and equivalents may be resorted to, falling within the scope of the present technology.

8.0 Concluding Thoughts on the Present Technology's Implications for Human Survival

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

[00638] 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 Sutskever (Chief Scientist, OpenAI), Bill Gates (Gates Ventures), Albert Efimov (Chief of
25 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
30 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.

[00639] 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?”

5 [00640] 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.

10 [00641] 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.

15 [00642] 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 do not know how the existing LLMs and AI agent systems they
20 have designed really work, let alone how to design AGI or SI, they also have no idea how to make these systems safe.

[00643] 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
25 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.

30 [00644] 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 which will allow the companies that adopt them to dominate the field of AI.

- [00645] The series of cited PPAs and PCTs, describe this “collective intelligence of intelligent (human and non-human) entities” approach to advanced AI, AGI, and SI systems. The present technology, 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.
- 5 [00646] 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
- 10 requirement for increasing intelligence. This data can come from any sufficiently intelligent entity, but may come from humans initially until non-human entities increase their level of intelligence.
- [00647] The present technology shows how to gather that data in a cost-effective and highly efficient manner. The present technology can generate huge amounts of revenue for online advertising companies, and even larger profits for companies that use the data to increase the
- 15 intelligence of their AI systems. The online Ad unit technologies 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 and ethical knowledge and preferences must be gathered from humans especially (and human-aligned intelligent entities more generally) as the present technology is rolled out.
- 20 [00648] 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 present technology for the benefit of humanity. If Alphabet fails, Meta, Amazon, Tencent, or other competitors might succeed. None of these companies want humanity to go extinct. What good are profits if no one is around to enjoy them?
- 25 [00649] The present technology provides a path to not only huge profits but also safe AGI.
- [00650] 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
- 30 agents. Rather, the best approach is for humans values and ethics to be inextricably intertwined with all of the knowledge that makes the AI systems intelligent in the first place.
- [00651] 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
5 (in the opinion of the humans) to create its solutions.

[00652] 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.

10 **[00653]** 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.

15 **[00654]** 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.

[00655] 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 are it.

CLAIMS

What is claimed is:

1. A system for utilizing online advertising technology for increasing an intelligence of an Artificial Intelligence (AI) agent or system, the system comprising:
 - 5 a computer system comprising: a processor, a computer-readable storage medium, and program instructions stored on the computer-readable storage medium being executable by the processor to cause the computer system to:
 - provide the online advertisement unit to one or more user computer systems utilizing a network;
 - 10 receive one or more inputs from the user computer systems participating in the online advertisement unit, the inputs being providing by an intelligent entity utilizing the user computer systems respective, and the inputs being associated with any one of or any combination of solving the problem, solving a sub-problem of the problem, and advancing progress on the problem, the intelligent entity being any one of or any
 - 15 combination of a human user utilizing a computer system, an AI agent or system, an Artificial General Intelligent (AGI) agent or system, and a SuperIntelligent (SI) agent or system;
 - communicate the inputs to the AI agent or system;
 - utilize the inputs by the AI agent or system in a universal problem solving method to
 - 20 generate a solution to the problem or the sub-problem; and
 - train the AI agent or system with results from the universal problem solving method that are based on successful or unsuccessful solution attempts to solve the problem or the sub-problem, thereby increasing an intelligence of the AI agent or system.
2. A method for online advertising technology utilized with an Artificial Intelligence (AI) agent or system for increasing an intelligence of the AI agent or system, the method comprising the steps of:
 - 25 receiving problem information of a problem provided by an AI agent or system;
 - populating an online advertisement unit including the problem information;
 - providing the online advertisement unit to one or more user computer systems utilizing a network;
 - 30 receiving one or more inputs from the user computer systems participating in the online advertisement unit, the inputs being providing by an intelligent entity utilizing the user computer systems respective, and the inputs being associated with any one of or any combination of solving the problem, solving a sub-problem of the problem, and advancing progress on the problem, the intelligent entity being any one of or any combination of a

- human user utilizing a computer system, an AI agent or system, an Artificial General Intelligent (AGI) agent or system, and a SuperIntelligent (SI) agent or system;
communicating the inputs to the AI agent or system;
utilizing the inputs by the AI agent or system in a universal problem solving method to generate
5 a solution to the problem or the sub-problem; and
training the AI agent or system with results from the universal problem solving method that are
based on successful or unsuccessful solution attempts to solve the problem or the sub-
problem, thereby increasing an intelligence of the AI agent or system.
3. The method of claim 2 further comprises the step of receiving advertisement specifications from a
10 client to be used in the populating or targeting of the online advertisement unit in combination with
the problem information.
4. The method of claim 3, wherein the advertisement specifications include any one of or any
combination of advertisement content, demographic information, location restrictions for the online
advertisement unit, advertisement budget, and metrics for determining a successful solving of the
15 problem or the sub-problem.
5. The method of claim 2 further comprises the step of billing the client based on cost per thousand
impressions or clickthrough rate metrics.
6. The method of claim 2, wherein the problem information includes any one of or any combination
of a current problem state of the problem, a sub-problem of the problem, a goal of the problem, a
20 sub-goal of the goal, an operator, and list of potential operators.
7. The method of claim 2 further comprises the step of communicating to the AI agent or system one
or more additional inputs from additional intelligent entities in combination with the inputs from the
participating intelligent entity, the intelligent entities being any one of or any combination of an
additional human user utilizing an additional computer system, an additional AI agent or system, an
25 additional Artificial General Intelligent (AGI) agent or system, and an additional SuperIntelligent
(SI) agent or system.
8. The method of claim 2, wherein the updating of the online advertisement unit includes any one of
or any combination of whether the inputs were accepted or rejected, whether a goal or sub-goal of
the problem has been achieved, if the input was selected by the AI agent or system to advance
30 problem solving, a description of a new problem state after applying the selected inputs, an amount
of credit or payment that the intelligent entities has accumulated based on inputs from the intelligent
entities, and an additional request for new user input by way of the online Ad unit or a link that
directs the intelligent entities to an interface for participating in the problem solving process on the
problem.

9. The method of claim 2 further comprises the step of crediting the intelligent entities if the problem or the sub-problem is solved, based on an amount of contribution by the participating intelligent entities.
10. The method of claim 2 further comprises the step of soliciting, if the problem or the sub-problem is not solved, additional input from the intelligent entity user of the participating user computer systems respectively.
11. The method of claim 2 further comprises the step of compensating the intelligent entity when the intelligent entity user exits the online Ad advertisement unit and stops participating in the problem solving process.
12. The method of claim 2, wherein the training of the AI agent or system further includes any one of or any combination of a training method selected from the group consisting of differential privacy, federated learning, homomorphic encryption, synthetic data generation, secure multi-party computation, data anonymization, transfer learning, active learning, self-supervised learning, domain adaptation, reinforcement learning, few-shot learning, explainable AI, privacy-preserving record linkage, data augmentation, generative adversarial networks, crowdsourcing for data labelling, model personalization layers, knowledge distillation, and ensemble learning.
13. The method of claim 2 further comprises the step of creating a database of intelligent entity experts, wherein the intelligent entity experts being any one of or any combination of human user utilizing a computer system, an additional AI agent or system, an Artificial General Intelligent (AGI) agent or system, and a SuperIntelligent (SI) agent or system.
14. The method of claim 13, wherein the online advertisement unit is provided to one or more of the intelligent entity experts from the database.
15. The method of claim 13, wherein the creating of the database utilizes any one of combination of a technology selected from the group consisting of relational database management system, NoSQL databases, data warehousing, data mining, machine learning algorithms, graph databases, vector databases, indexing, full-text search, blockchain, data visualization tools, application programming interfaces, data cleaning and preprocessing, cloud storage, caching, transactional database systems, real-time database systems, content delivery network, data compression, data encryption, and replication.
16. The method of claim 2 further comprises the step of bidding on attention using an attention spot market.
17. The method of claim 16, wherein the attention spot market includes:
a means for the intelligent entity with attention to sell to access a marketplace and specify seller information for sale and an ask price for the information;

- a means for buyers of intelligent entity attention to access the marketplace and specify buyer information the buyer is willing to buy and a bid price for buying the buyer information;
a market mechanism for queuing the bid price and the ask price, including categories of the buyer and seller information, wherein each category has a market in the marketplace;
5 the market mechanism is configured or configurable to make the market in each category by market makers; and
the market mechanism is configured or configurable to match bid prices and ask prices, and a transaction occurs that is binding on the buyer and the seller of the information.
18. The method of claim 16, wherein the attention spot market includes the steps of:
- 10 buyers and sellers of intelligent entity attention and expertise register on a platform, the buyers providing details about buyer interests or expertise, the sellers providing details about sellers interests or expertise;
listing by the sellers available time slots and expertise areas, and listing by the buyers needs and time slots the buyers is interested in;
15 utilizing, by the platform, an algorithm to dynamically price intelligent entity attention and expertise based on supply, demand, and user ratings;
matching one or more of the buyers and one or more of the sellers based on requirements, availability, and price;
enabling transactions where the buyers pay for the seller time slots, and wherein the platform
20 takes a commission; and
providing feedback where after each session, the buyers and the sellers rate each other, influencing future pricing and matching.
19. The method of claim 18, wherein the platform utilizes one or a combination of mechanisms selected from the group consisting of Dutch auction, reverse auction, sealed bid auction, open
25 ascending price auction, fixed price with time priority, dynamic pricing based on ratings, supply-demand curve adjustment, time slot segmentation, subscription access, freemium model, group buying, tiered expertise levels, flash sales, loyalty points system, hybrid auction, geographic pricing, behavioral pricing, escrow system, social influence pricing, and tokenized transactions.
20. The method of claim 16, wherein the attention spot market includes the steps of:
- 30 buyers and sellers of intelligent entity attention and expertise register on a platform, the buyers providing details about buyer interests or expertise, the sellers providing details about sellers interests or expertise;
creating an auction by the sellers for seller time slots;
placing one or more buyer bids by a buyer on time slots and expertise the buyer requires;

closing the auction at a predetermined time or when the seller accepts a buyer bid;
paying the seller by the buyer of a winning bid, and the seller provides the attention or expertise
based on the time slot, wherein the platform mediates an exchange and secures payment; and
providing feedback where after each auction, the buyers and the sellers rate each other, affecting
5 future auctions and visibility on the platform.

21. The method of claim 20, wherein the platform utilizes one or a combination of mechanisms
selected from the group consisting of Dutch auction, reverse auction, sealed bid auction, open
ascending price auction, fixed price with time priority, dynamic pricing based on ratings, supply-
demand curve adjustment, time slot segmentation, subscription access, freemium model, group
10 buying, tiered expertise levels, flash sales, loyalty points system, hybrid auction, geographic pricing,
behavioral pricing, escrow system, social influence pricing, and tokenized transactions.

22. The method of claim 2, wherein the online advertisement unit is configured or configurable to
include material relevant to an area of expertise associated with the problem information.

23. The method of claim 2, wherein the online advertisement unit is configured or configurable to,
15 upon interaction by the intelligent entity, would expand or redirect the intelligent entity to a secure
form on a landing page, wherein the form is configured to collect user information and to provide
questions to the intelligent entity regarding to qualifications of the intelligent entity.

24. The method of claim 23 further comprising the step of validating the user information and
categorizing the intelligent entity in a database based on the qualifications.

20 25. The method of claim 2, wherein the online advertisement unit is configured or configurable to
include an interface allowing the intelligent entity to input a solution to the problem or the sub-
problem.

26. The method of claim 25, wherein the online advertisement unit is configured or configurable to
include any one of or any combination of elements of gamification, a scoring system, and a
25 leaderboard of other intelligent entities.

27. The method of claim 25, wherein the online advertisement unit is configured or configurable to
include text input fields or other interactive tools enabling the intelligent entity to provide a solution
to the problem or sub-problem directly within a space of the online advertisement unit.

28. The method of claim 27, wherein the solution is stored in a database, and an algorithm is utilized
30 to assess a quality and relevance of the solution based on predefined criteria.

29. The method of claim 28, wherein the intelligent entity receive instant feedback or rewards points
based on the assessed quality and relevance of the solution, redeemable for various incentives.

30. The method of claim 28, wherein contributions by the intelligent entity are tracked to identify
top contributors from multiple additional intelligent entity for potential future engagement.

31. The method of claim 25, wherein the online advertisement unit is configured or configurable to, upon interaction by the intelligent entity, would expand or redirect the intelligent entity to a secure form on a landing page, wherein the form is configured to collect user information and to provide questions to the intelligent entity regarding to qualifications of the intelligent entity.
- 5 32. The method of claim 31 further comprising the step of validating the user information and categorizing the intelligent entity in the database based on the qualifications.
33. The method of claim 2 further comprises the steps of:
- analyzing the inputs from the participating intelligent entity to determine a quality or relevance of the inputs, respectively; and
- 10 compensating the participating intelligent entity based on the quality or relevance of the inputs.
34. The method of claim 33 further comprises the step of providing digital wallet information by each of the participating intelligent entity into a secure form, wherein the secure form is within the online advertisement unit or is provided by a linked platform.
35. The method of claim 33, wherein the compensation is provided in any one of or any
- 15 combination of cryptocurrency, electronic gift cards, electronic vouchers, subscriptions, access to online resources or content, recognition on online platforms, merchandise related to an expertise of the intelligent entity respectively, and sponsorships for professional conferences or events.
36. The method of claim 33 further comprises the step of incentivizing the participating intelligent entity to refer other intelligent entity to an online platform, by offering additional compensation for
- 20 each referral that contributes to the platform, wherein the online platform receives the problem and provides the online advertisement unit.
37. The method of claim 2 further comprising the step of receiving reputational metrics on each of the intelligent entity.
38. The method of claim 37, wherein the reputational metrics are multi-dimensional including any
- 25 one of or any combination of timeliness, budget compliance, quality of work, solution success rate, client satisfaction, peer-rating of competency, external reputation, innovation score, communication skills, adaptability, leadership quality, technical proficiency, learning agility, conflict resolution, project management skills, reliability, efficiency, cultural fit, work ethic, client retention rate, feedback responsiveness, networking ability, mentorship and training, availability and
- 30 responsiveness, problem-solving speed, and creativity and innovation.
39. The method of claim 37, wherein the reputational metrics for any one of the intelligent entity are automatically updated within the online advertisement unit.

40. The method of claim 2 further comprising the step of updating the online advertisement unit based on the results from the universal problem solving method conducted on the problem or the sub-problem.

5 41. The method of claim 2, wherein the online advertisement unit is configured or configurable to ask the intelligent entity to provide a description of how the problem is to be framed, and then the description is then translated into a language of the universal problem solving method using a natural language to problem solving language translator.

42. The method of claim 2, wherein the online advertisement unit is configured or configurable to ask the intelligent entity to provide a description of intermediate goals towards a final solution of the
10 problem, and wherein the intermediate goals are communicated to the AI agent or system for processing by the universal problem solving method.

43. The method of claim 2 further comprises the step of requesting, by the AI agent or system, the intelligent entity to provide safety or ethical information associated with the problem, the sub-problem or goals associated with the problem.

15 44. The method of claim 43 further comprises the step of utilizing the safety or ethical information in the populating or in updating the online advertisement unit.

45. The method of claim 2 further comprises the step of modifying the online advertisement unit to ask the intelligent entity to vote, rank or rate one or more potential goals, sub-goals or actions that have been proposed by any one of or any combination of the intelligent entity, and intelligent
20 entities.

46. The method of claim 2 further comprises the step of requiring the AI agent or system to be identified as an AI agent so any inputs from the AI agent or system is excluded where only human opinions are sought.

47. The method of claim 2, wherein the online advertisement unit is provided by an online
25 advertisement service provider that uses user information about the intelligent entity to target the intelligent entity with expertise relevant to the problem.

48. The method of claim 47 further comprises the step of measuring and recording response rates for the participating intelligent entity who contribute information within the online advertisement unit or by way of links to websites or interfaces outside the online advertisement unit.

30 49. The method of claim 47 further comprises the step of recording metrics related to a quality of information or solutions provided by the targeted intelligent entity.

50. The method of claim 49 further comprises the step of calculating correlations or statistical relationships on the user information to determine any one of or any combination of which factors have a highest impact on the metrics, which of the metrics is prioritized by the online advertisement

service provider, and which of the inputs was most effective at meeting or maximizing problem solving or training of the AI agent or system.

51. The method of claim 47 further comprises performing arbitrage including the steps of:

5 estimating, by the online advertisement service provider, a value of a particular type of intelligent entity attention and knowledge with regard to the problem or sub-problem;
estimating, by the online advertisement service provider, a cost of obtaining an amount of the intelligent entity attention required to complete the problem or sub-problem from the intelligent entity or intelligent entities by way of the online advertisement unit; and
10 determining if the estimated value exceeds the estimated cost by a predetermined or dynamic variable representing a profit margin, then purchase the intelligent entity attention at cost approximating that of the estimated value and sell the intelligent entity attention to clients, or use the intelligent entity attention to perform tasks that create value approximating that of the estimated cost.

52. The method of claim 51 further comprises the step of performing arbitrage in priority order of a
15 largest arbitrage opportunities first, until a minimal acceptable arbitrage opportunity is reached.

53. The method of claim 2, wherein the online advertisement unit includes one or more interfaces selected from the group consisting of Large or Small Language Model interfaces, natural language text, audio-based interfaces, text boxes, dropdown lists, templates, dynamically sizing input and output areas, visual input or output devices, virtual reality devices, and augmented reality devices.

20 54. The method of claim 2, wherein the overall online advertisement unit is comprised of multiple online advertisement units with a first online advertisement unit being provided to a first set of the intelligent entities for receipt of the inputs, and a second online advertisement unit being provided to the first set or a second set of the intelligent entities and asking to vote on the inputs.

25 55. The method of claim 2, wherein the inputs are provided to the AI agent or system in real-time for real-time utilization by the universal problem solving method.

56. A method of increasing monetization of online advertising revenue for an online advertisement service provider by delivering to a client information based on intelligent entity attention, the method comprising the steps of:

30 building, by an online advertisement service provider, one or more databases containing information about intelligent entities, the information is related to targeting specific types of online advertisements to specific intelligent entities, the intelligent entities being any one of or any combination of a human user utilizing a computer system, an AI agent or system, an Artificial General Intelligent (AGI) agent or system, and a SuperIntelligent (SI) agent or system;

communicating to the online advertisement service provider by a client a request to purchase intelligent entity attention, the request including requirements and desired categories of the intelligent entity attention;

performing checks on the request by the online advertisement service provider at a time that the requirements are communicated, the checks are configured or configurable to ensure that the request does not violate regulations or ethical requirements;

creating content for an interactive online advertisement unit that is configured or configurable to capture intelligent entity attention and work from intelligent entities;

deploying, by the online advertisement service provider, the interactive online advertisement unit to a computer system of the intelligent entities by way of a network;

allowing the interactive online advertisement unit to collaborate with multiple of the intelligent entities and integrate with real-time or asynchronous data capabilities;

performing ethics and safety checks during a problem solving process on a problem provided by the client, the ethics and safety checks are configured or configurable to ensure that unethical or unsafe expertise is not part of the problem; and

improving the interactive online advertisement unit based on a feedback loop including on any one of or any combination of reputational metrics, and metrics related to the problem solving process or knowledge captured.

57. The method of claim 56, wherein the client is any one of or any combination of an Artificial Intelligent (AI) agent or system, a system that trains or customizes AI agents, and a system that provides solutions to problems.

58. The method of claim 56 further comprises the step of determining if the database does not include a predetermined number of the intelligent entities, or if additional intelligent entities are required, then configuring the interactive online advertisement unit to acquire additional intelligent entities.

59. The method of claim 56, wherein the requirements are communicated by way of interaction with an attention spot market.

60. The method of claim 59, wherein the attention spot market includes any one of or any combination of:

a means for the intelligent entities with attention to sell to access a marketplace and specify seller information for sale and an ask price for the information;

a means for buyers of intelligent entity attention to access the marketplace and specify buyer information the buyer is willing to buy and a bid price for buying the buyer information;

a market mechanism for queuing the bid price and the ask price, including categories of the buyer and seller information, wherein each category has a market in the marketplace; the market mechanism is configured or configurable to make the market in each category by market makers; and

- 5 the market mechanism is configured or configurable to match bid prices and ask prices, and a transaction occurs that is binding on the buyer and the seller of the information.

61. The method of claim 59, wherein the attention spot market includes the steps of:

- buyers and sellers of intelligent entity attention and expertise register on a platform, the buyers providing details about buyer interests or expertise, the sellers providing details about sellers
10 interests or expertise;

listing by the sellers available time slots and expertise areas, and listing by the buyers needs and time slots the buyers is interested in;

utilizing, by the platform, an algorithm to dynamically price intelligent entity attention and expertise based on supply, demand, and user ratings;

- 15 matching one or more of the buyers and one or more of the sellers based on requirements, availability, and price;

enabling transactions where the buyers pay for the seller time slots, and wherein the platform takes a commission; and

- 20 providing feedback where after each session, the buyers and the sellers rate each other, influencing future pricing and matching.

62. The method of claim 61, wherein the platform utilizes one or a combination of mechanisms selected from the group consisting of Dutch auction, reverse auction, sealed bid auction, open ascending price auction, fixed price with time priority, dynamic pricing based on ratings, supply-demand curve adjustment, time slot segmentation, subscription access, freemium model, group
25 buying, tiered expertise levels, flash sales, loyalty points system, hybrid auction, geographic pricing, behavioral pricing, escrow system, social influence pricing, and tokenized transactions.

63. The method of claim 59, wherein the attention spot market includes the steps of:

- buyers and sellers of intelligent entity attention and expertise register on a platform, the buyers providing details about buyer interests or expertise, the sellers providing details about sellers
30 interests or expertise;

creating an auction by the sellers for seller time slots;

placing one or more buyer bids by a buyer on time slots and expertise the buyer requires;

closing the auction at a predetermined time or when the seller accepts a buyer bid;

paying the seller by the buyer of a winning bid, and the seller provides the attention or expertise based on the time slot, wherein the platform mediates an exchange and secures payment; and providing feedback where after each auction, the buyers and the sellers rate each other, affecting future auctions and visibility on the platform.

- 5 64. The method of claim 63, wherein the platform utilizes a mechanism selected from the group consisting of Dutch auction, reverse auction, sealed bid auction, open ascending price auction, fixed price with time priority, dynamic pricing based on ratings, supply-demand curve adjustment, time slot segmentation, subscription access, freemium model, group buying, tiered expertise levels, flash sales, loyalty points system, hybrid auction, geographic pricing, behavioral pricing, escrow system, 10 social influence pricing, and tokenized transactions.

65. The method of claim 59 further comprises the step of providing a feedback mechanism associated with the intelligent entity attention spot market, the feedback mechanism being configured or configurable to record each and every step in the problem solving process, and to create a vector track record of a performance of all the intelligent entities on every task and sub-task.

- 15 66. The method of claim 65, wherein the vector track record is implemented by way of blockchain technology to allow for precise reputations that are analyzed by an AI agent or system and converted into estimates of a value for each of the intelligent entities for any of the tasks.

67. The method of claim 56, wherein the interactive online advertisement unit is configured or configurable to capture the intelligent entity attention and work from within the interactive online advertisement unit or by providing link to one or more webpages or interfaces that are optimized for 20 the capture outside of the interactive online advertisement unit.

68. The method of claim 56, wherein the building of the databases includes the step of purchasing, by the online advertisement service provider, user data and information from remote sources.

- 25 69. A method for online advertising technology utilized with an Artificial Intelligence (AI) agent or system, the method comprising the steps of:

populating, by an AI agent or system, an online advertisement unit including problem information provided to the AI agent or system;

providing the online advertisement unit to one or more user computer systems utilizing a network;

- 30 receiving one or more inputs from the user computer systems participating in the online advertisement unit, the inputs being providing by an intelligent entity respectively, and the inputs being associated with any one of or any combination of solving the problem, solving a sub-problem of the problem, and advancing progress on the problem, the intelligent entity being any one of or any combination of a human user utilizing a computer system, an AI

- agent or system, an Artificial General Intelligent (AGI) agent or system, and a SuperIntelligent (SI) agent or system;
- communicating the inputs to the AI agent or system;
- utilizing the inputs by the AI agent or system or an additional AI agent or system in a universal
- 5 problem solving method to generate a solution to the problem or the sub-problem; and
- utilizing a feedback mechanism that provides metrics to the AI agent or system to train the AI agent or system, the metrics assign credit or blame to the participating intelligent entity and targeting mechanisms that directs the online advertisement unit to particular intelligent entities.
- 10 70. The method of claim 69 further comprises the step of measuring and recording response rates for the participating intelligent entities who contribute information within the online advertisement unit or by way of links to websites or interfaces outside the online advertisement unit.
71. The method of claim 69 further comprises the step of recording the metrics related to a quality of information or solutions provided by the targeted intelligent entities.
- 15 72. The method of claim 71 further comprises the step of calculating correlations or statistical relationships on the user information to determine any one of or any combination of which factors have a highest impact on the metrics, which of the metrics is prioritized by the online advertisement service provider, and which of the inputs was most effective at meeting or maximizing problem solving or training of the AI agent or system.
- 20 73. The method of claim 69 further comprises performing arbitrage including the steps of:
- estimating, by the online advertisement service provider, a value of a particular type of intelligent entity attention and knowledge with regard to the problem or sub-problem;
- estimating, by the online advertisement service provider, a cost of obtaining an amount of the intelligent entity attention required to complete the problem or sub-problem from the
- 25 intelligent entity or intelligent entities by way of the online advertisement unit; and
- determining if the estimated value exceeds the estimated cost by a predetermined or dynamic variable representing a profit margin, then purchase the intelligent entity attention at cost approximating that of the estimated value and sell the intelligent entity attention to clients, or use the intelligent entity attention to perform tasks that create value approximating that of
- 30 the estimated cost.
74. The method of claim 73 further comprises the step of performing arbitrage in priority order of a largest arbitrage opportunities first, until a minimal acceptable arbitrage opportunity is reached.