



Agents of exploration and discovery

Pat Langley

Institute for Defense Analyses

Correspondence

Pat Langley, Information Technology and Systems Division, Institute for Defense Analyses, 4850 Mark Center Drive, Alexandria, VA 22311 USA.
Email: patrick.w.langley@gmail.com

Abstract

Autonomous agents have many applications in familiar situations, but they also have great potential to help us understand novel settings. In this paper, I propose a new challenge for the AI research community: developing embodied systems that not only explore new environments but also characterize them in scientific terms. Illustrative examples include autonomous rovers on planetary surfaces and unmanned vehicles on undersea missions. I review two relevant paradigms: robotic agents that explore unknown areas and computational systems that discover scientific models. In each case, I specify the problem, identify component functions, describe current abilities, and note remaining limitations. Finally, I discuss obstacles that the community must overcome before it can develop integrated agents of exploration and discovery.

BACKGROUND AND MOTIVATION

Recent years have seen substantial progress and corresponding excitement about autonomous agents for human environments, such as self-driving vehicles and delivery drones. However, there is also a need for agents that operate in unfamiliar, remote settings to pursue more open-ended research activities. Examples include robots that explore the surfaces of other planets, traverse deep ocean trenches, and navigate subterranean caverns to understand environments that differ radically from our own. We need artifacts that follow in the footsteps of early scientific explorers like Charles Darwin and Alfred Wallace, who not only collected samples on their extended voyages but who also noted relations among them and formed suggestive hypotheses about the processes that produced them.

I will refer to such systems as *agents of exploration and discovery*, as they combine two facets of human intelligence that have previously been studied separately. First is the ability to explore a novel environment physically, visiting locations and mapping their features. Second is the capacity to pursue scientific research in this context, collecting measurements, finding regularities, forming hypotheses, and gathering additional data to test them. There has been some work along these lines with plane-

tary rovers and unmanned underwater vehicles, but they have mainly been teleoperated and they have focused on collection of data rather than on their interpretation. In the longer term, we will need truly autonomous explorer scientists that, like the one envisioned in Zelazny's (1975) *Home is the Hangman*, visit environments humans cannot tolerate and communicate discoveries to their developers in familiar and understandable terms.

The next section presents some motivating examples of such intelligent systems. I then examine the two core abilities – exploring novel environments and characterizing them scientifically – in some detail. In each case, I define the task addressed, identify component abilities, cite relevant work, and note the state of the art. After this, I discuss integrated agents that combine these two capacities, posing a number of open challenges that could drive future research in this area. The closing section reviews the key ideas and encourages readers to tackle this important and audacious problem.

ILLUSTRATIVE EXAMPLES

Consider a scenario in which a planetary rover lands on Venus to explore and understand its little-known surface.



The planet's dense cloud cover means that, unlike Mars and the Moon, detailed satellite atlases are unavailable. The rover must explore the region and build informative maps while avoiding local hazards. High temperature and pressure make the surface an unlikely host for life, but the geology and weather hold great scientific interest. To this end, the rover would identify categories of rock formations and weather patterns, induce laws that govern their behavior, and infer mechanisms (like erosion and tectonics) that may have produced them. Much of the data would come from passive observation, with the rover selecting promising locations for measurements, but it could also carry out focused experiments to determine which chemical reactions are supported by the environment. The distance and time lag would make communication with Earth difficult, so both exploration and discovery would be primarily autonomous.

Another scenario involves an unmanned underwater vehicle on a mission to observe and understand phenomena on the ocean floor, including hydrothermal vents. The great depth means that only coarse maps will be available in advance, so the agent must visit sites of interest to create detailed descriptions. These will include information about temperature, pressure, salinity, and currents, as well as geological structures and seafloor covers. Using passive sensors, the vehicle would estimate distributions of species, relating them to conditions like the presence of heat sources and nutrients. The agent would identify food webs in different areas and their connection to geological and ocean processes like volcanic activity. Another scientific target would be chemical and biochemical reactions enabled by extreme pressures and temperatures on the ocean floor. Again, communication would be very limited, so the agent would operate autonomously for extended periods, carrying out much of its data analysis on board.

This does not limit the settings in which autonomous agents could combine exploration with discovery. The integrated technology would be relevant to any remote physical environment in which humans will find it difficult to survive, much less pursue systematic research. These include equatorial jungles, polar regions, subterranean caverns, deep ocean trenches, dense forest canopies, active volcanic slopes, and airless asteroids. In such challenging surroundings, robotic explorer scientists can serve as proxies that carry out extended and comprehensive research missions.¹

AGENTS OF EXPLORATION

Encounters with entirely new milieus have often led to scientific advances, but they are invariably associated with careful descriptions of those environs. Because the

autonomous agents envisioned will operate in unfamiliar settings, they must explore their locales in order to find content of scientific interest. We can specify this exploration task as:

- *Given*: An environment to be traversed and described;
- *Given*: A mobile robot equipped with appropriate sensors and effectors;
- *Given*: Procedures and criteria for using these sensors and effectors; and
- *Find*: A descriptive map of the environment with features of interest.

This statement is abstract, but our aim is to provide a framework for discussion rather than any specific solution. For instance, the generated map may be a simple grid or a more structured description with distinct places and regions. Each location on the map may be annotated with qualitative content (e.g., water or land), quantitative attributes (e.g., elevation, salinity), or other descriptors. Robotic map construction is a standard problem that has been studied for decades, but we should review which facets are relevant to our agenda.

Facets of robotic exploration

Research on robotic exploration of large-scale environments has a long history, but some distinctions in the field are especially important to extended, remote missions:

- *Teleoperation versus autonomy*. Early mobile robots were teleoperated by humans, and this continues for many applications, including most instances of planetary exploration. Robotic autonomy has since become far more common, but it occurs mainly in highly structured settings, such as cars that drive themselves on city streets and highways. Exploration of truly remote sites will require more nearly autonomous agents that are out of touch and on their own for extended periods.
- *Types of effectors*. Mobile robots have been developed for a variety of environmental conditions. In addition to wheeled vehicles for driving, there has been work on legged movement over natural terrains, submersible vehicles for underwater traversal, and flying platforms for aerial operation. Different forms of locomotion will be appropriate for different types of environments (e.g., caverns vs. underwater trenches), but most alternatives will support the types of exploration that we desire.
- *Types of sensors*. The robotics community has created many devices for observing the agent's environment, including sonar, lasers, LiDAR, stereo vision, and other



range finders that let agents infer their surrounding structures. These may be complemented by standard sensors for temperature, pressure, and other ambient variables, as well as more sophisticated devices like mass spectrometers. Different devices will be useful for distinct environments and scientific objectives.

Of course, researchers have also developed numerous robotic architectures that combine sensors and effectors with internal processing to infer descriptions of agent surroundings and to respond accordingly. I will not delve into issues of architectural design here, as many alternative frameworks can support the functionality that is needed for our purposes.

However, I should review the constituent abilities that are required for effective robot exploration as revealed by previous work on the topic. These capacities include:

- *Localization*, which determines the agent's position and orientation on a local or large-scale map. This process involves comparing current sensor readings to those predicted by different hypotheses, ranking candidates by their degree of match, and selecting one or more options from this set. In continuous localization, the robot must update its position and orientation estimates repeatedly, as the platform traverses the environment and collects new sensor readings.
- *Target selection*, which chooses one or more destinations that the robot should visit during its exploration activities. This process involves evaluating candidate locations in terms of the agent's objectives, selecting a subset of them based on these criteria, and deciding on their relative priorities. Decisions about targets can change as the robot moves through an area and as additional information becomes available, suggesting new candidates and changing the scores for old ones.
- *Navigation*, which generates a route that the agent can follow from its current location to one or more selected targets. This process includes both high-level route planning, which often requires search through a space of possible paths, and low-level obstacle avoidance, which typically draws on reactive control. Decisions to change target locations can force the agent to revise its route plans, while dynamic environments can require rapid responses at the control level.
- *Mapping*, which extends or refines the agent's current map based on sensor readings that it collects during traversal of planned routes or during visits to target destinations. This process can range from low-level updates about the occupancy of grid cells to high-level structural changes about places and spatial relations among them. Effective map construction relies on accurate localiza-

tion of the robot for proper estimates about its position and orientation.

Of course, mobile robots must also deal with maintenance issues, such as ensuring they have enough power to continue their missions, but the four elements described above form the core mechanisms of exploration. More generally, such agents must address the tradeoff between exploration and exploitation, although the former will have higher priority in our scenarios than in many classic robotics applications.

State of robotic exploration

One might naturally ask whether these technical solutions are mature enough for robotic exploration on extended, autonomous missions to distant environments. Fortunately, 50 years of research and development have produced a suite of robust representations and mechanisms that have been tested extensively not only in laboratories but also in many less constrained outdoor settings. The component tasks are well defined, algorithms for addressing them are well understood, and integrated robotic systems have combined them in effective ways. Textbooks (e.g., Stachniss 2009; Sturges 2015) and edited collections (e.g., Ishigami and Yoshida 2021) have documented many of the advances in this arena, some of them computational in character and others due to new sensors and effectors.

The literature on robotic exploration is broad and extensive. Early work on the acquisition and use of maps focused on this task in isolation. Probabilistic occupancy grids emerged as a key representation for map construction (Moravec and Elfes 1985), but they have also proved useful for localization (Yamauchi and Langley 1997). Later efforts combined these ideas into methods for simultaneous localization and mapping (Yamauchi, Schultz, and Adams 1998; Durrant-Whyte and Bailey 2006), which have been strengthened by techniques like particle filtering (Aru-lampalam et al. 2002) to deal with position uncertainty and by combining topological and metric approaches (Besson, Modayil, and Kuipers 2010) to handle large spaces. Robotic platforms can now acquire accurate, large-scale maps quite rapidly, and early work on indoor exploration has been superseded by systems that operate in rich outdoor environments (Thrun and Montemerlo 2005). These and related techniques have been used in many different physical settings, including self-driving cars (Bresson et al. 2017), underwater vehicles (Hidalgo and Braunl 2015; Zhao et al. 2019), aerial drones (Li, Savkin, and Vucetic 2020), underground robots (Ebadi et al. 2020), and even planetary rovers (Geromichalos et al. 2020). The last four topics are

especially relevant for agents that must pursue extended remote missions.

Despite this progress, research on robotic exploration has two drawbacks that limit its usefulness for scientific aims. The first is that the constructed maps emphasize voxel-level summaries of the environment rather than more complex descriptions. These maps can incorporate values for multiple variables and even detected objects, but they generally do not move beyond annotated grids to higher-level scientific accounts. There has been some progress in “semantic mapping” (Kostavelis and Gasteratos 2015; Rosinol et al. 2020) that assigns labels to locations and infers spatial structures, but this area is far less advanced than basic map construction and needs further development. The second is that current robotic platforms depend on humans to specify many aspects of their missions, as they lack the ability to pursue open-ended scientific exploration. The field needs additional work on generating and selecting high-level tasks, possibly borrowing from the recent literature on goal reasoning (Aha et al. 2013). We must extend the current exploration paradigm to overcome both of these limitations before it can help turn our vision into reality.

AGENTS OF DISCOVERY

Although exploratory agents are essential to our vision, creating detailed maps alone will not be sufficient. We must combine robotic explorers with agents of discovery that note and explain regularities in the environments they traverse. We can specify this task more explicitly as:

- *Given*: An environment to be understood scientifically;
- *Given*: High-level scientific objectives of the mission;
- *Given*: Observations taken from this environment;
- *Find*: Hypotheses and models stated in established scientific formalisms; and
- *Find*: New observations relevant to testing these claims.

This statement is similar to those for machine learning and data mining, but differs in its emphasis on generating models in scientific formalisms (Džeroski, Langley, and Todorovski 2007), rather than notations invented by computational researchers. These differ for different fields – equations in physics, reaction pathways in chemistry, food webs in ecology – but they ensure that findings are accessible to human scientists.

Facets of scientific discovery

Historians of science study many different aspects of the research process, but they make three high-level distinc-

tions among activities that are especially important for devising agents of discovery:

- *Observation versus experimentation*. Some scientific activities focus on the collection and analysis of samples available in the environment, while others emphasize controlled studies that determine some variables’ effects on others.² Certain fields, such as astronomy, geology, and meteorology, commonly rely on observational studies because they often lack experimental control. Others, such as physics, chemistry, and biology, began as primarily observational but have shifted almost entirely to experimentation as they have matured. The exploratory character of the envisioned agents means they will be largely observational, although they must still decide which data to collect and where to take measurements.
- *Qualitative versus quantitative accounts*. Although mature scientific fields are usually associated with numeric laws and models, they are typically preceded by the discovery of qualitative relations that provide context for them. This ordering may seem counterintuitive to readers who are familiar with recent results in statistical learning, which often cast problems in terms of regression. However, the precedence of qualitative over quantitative discovery has been well documented in the history of science (Mason 1962). Chemists detect the existence of particular reactions before they determine details about their proportions and rates, and ecologists establish food webs before they identify numeric functions that control populations. Our agents should engage in both forms of scientific inquiry, but they will often need to develop qualitative accounts before they attempt to quantify them.
- *Normal versus revolutionary science*. Many popular accounts of discovery focus on rare breakthroughs like universal gravitation or evolution, but typical activities involve what Kuhn (1962) has called *normal science*. Here, researchers adopt shared assumptions and goals, and they aim at filling in a framework’s details rather than replacing it. Such paradigmatic frameworks provide concepts and constraints that serve to canalize the scientific process, making it tractable and enabling incremental progress. For this reason, agents that operate in novel environs should favor normal science, relying on established theories of hydrology, chemistry, and ecology to guide their discoveries, although some findings (e.g., extremophiles near thermal vents) may border on revolutionary.

Discovery systems should also mimic the incremental and cumulative character of human science. With a few exceptions like astronomy, the latter seldom acquires and processes samples in large batches. Rather, it accumulates



observations a few at a time, then uses them to extend and revise existing laws and models. Our agents of discovery should follow a similar strategy.

The history of science also shows that the discovery process occurs in identifiable stages, with later forms typically building on earlier ones (Mason 1962). These steps include:

- *Taxonomy formation*, which organizes entities with similar characteristics into classes and subclasses, often in a hierarchy. For example, chemistry identifies different types of substances, geology specifies rock categories, and biology distinguishes among life forms. Such ontologies define the basic terms that can appear in other scientific statements. They are generally stable, but they are extended and occasionally restructured in response to new observations.
- *Descriptive law induction*, which infers qualitative relations among classes of entities or numeric equations that connect variables. These summarize regularities in the environment using terms from the taxonomy. For instance, chemistry specifies which substances react to form others, but also states the relative masses involved in each reaction. Similarly, the specific heat law declares not only that two objects will approach equilibrium but also how final temperature relates to initial temperatures, mass, and specific heat.
- *Explanatory model construction*, which posits unobserved structures or processes that, if present, would account for observations or laws that summarize them. These hypothesize the underlying reasons for observed regularities, typically stated in terms of familiar concepts.³ For example, the atomic theory asserts that physical substances comprise tiny molecules, with different types made from different elements. The germ theory states that diseases result from microorganisms that spread throughout the population.

Autonomous discovery agents would engage in all of these scientific activities, each of which can involve detecting and responding to anomalous observations. They would identify new types of rock formations and organisms, form generalizations about when these entities occur and how they interact, and generate explanations of these phenomena in terms of structures and processes consistent with existing physical, chemical, geological, and biological theories.

State of computational scientific discovery

Some readers may question whether our understanding of the discovery process is mature enough for autonomous application to new environments. However, over 40 years

of research has produced a substantial literature on computational discovery using heuristic search, guided by data, through a space of models stated in established scientific formalisms. Edited volumes and special issues by Shrager and Langley (1990), Sleeman, Corruble, and Valdés-Pérez (2000), Džeroski and Todorovski (2007), and Addis et al. (2019) have documented intellectual advances, and Langley (2000) has reviewed efforts that produced publications in refereed journals devoted to a variety of natural sciences.

Early research in the area focused on induction of numeric equations (Langley 1981; Langley and Zytkow 1989), sometimes referred to as *equation discovery* (Todorovski 2011). Initial results involved examples from the history of physics and chemistry (Langley et al. 1987), but later endeavors have produced novel results in metallurgy (Mitchell et al. 1997), ecology (Todorovski, Džeroski, and Kompore 1998; Todorovski et al. 2003), and other scientific fields. Work on the problem has continued unabated, with more recent efforts examining classical physics (Wu and Tegmark 2019), fluid dynamics (Bruntona, Proctor, and Kutz 2016; Raissi and Karniadakis 2018), visual attention (Lane et al. 2016), and macroeconomics (Georgiev and Kazakov 2015). A parallel line of research has focused on inducing qualitative relations, leading to new results in medicine (King and Srinivasan 1996), biochemistry (Finn et al. 1998), number theory (Colton, Bundy, and Walsh 2000), and other disciplines. Unlike many approaches to machine induction, these systems do not require large training sets to produce useful results.

Analysis of techniques for discovery of descriptive laws clarify why the field has made such progress and why it may be ready for use in autonomous missions. As Langley, Shrager, and Saito (2002), among others, have argued, we can always specify a space of possible descriptive laws and search this space to find relations that fit observations. This search need not be blind, but can be guided by heuristics that combine regularities in the data, known domain constraints, and high-level criteria like parsimony. This lets one automate search of large spaces of hypotheses that cannot be traversed exhaustively. Naturally, this approach depends on some framework to define the space, but we are concerned here with normal science that operates within an established paradigm. When controlled experimentation is possible, it provides additional power by decomposing the discovery task into subtasks to further aid tractability (Langley 1981). Similar techniques have produced many successes in the data-mining community; the key difference lies in the use of established scientific formalisms (e.g., numeric equations) to ensure that results are interpretable.

Another line of research has addressed the construction of deeper models that explain phenomena in terms



of hypothesized structures or processes. Some early results dealt with historical reconstructions from chemistry (Langley et al. 1987), biochemistry (Kulkarni and Simon 1990), and particle physics (Kocabas 1991). However, predating these efforts was the DENDRAL project (Lindsay et al. 1980), which constructed structural models of molecules to explain mass spectroscopy measurements and which led to many publications in the organic chemistry literature.⁴ Later research in a similar spirit has produced refereed papers that reported new findings about reaction pathways in catalytic chemistry (Valdés-Pérez 1994) and metabolic regulation (King et al. 2004; King, Rowland, and Oliver 2009). More recently, the ACE system (Anderson et al. 2014), which infers process models that explain rock nuclide densities, has been widely used by geologists. Some work has combined the two themes by embedding equation discovery within construction of explanatory models (e.g., Bridewell et al. 2008). This hybrid approach has been applied to aquatic ecosystems (Atanasova et al. 2008) and epidemiology (Tanevski, Todorovski, and Džeroski 2016). The notion of explanation here is quite different from that adopted by recent AI research on the behavior of learned classifiers or policies. The aim is not to elucidate internal decisions but rather to explain external phenomena in terms of familiar concepts (Langley 2020).

These successes suggest that computational discovery of explanatory models also holds promise for use in remote settings, where it could propose accounts of observed phenomena that move beyond descriptive summaries. Like mechanisms for law discovery, methods for constructing explanatory models carry out constrained search through a space of candidates. The main difference is their introduction of theoretical entities or processes from which one can derive a set of observed phenomena, which may themselves be descriptive laws. Such models invariably draw on background knowledge and thus fall within an established paradigm, although this knowledge is often generic and thus requires little domain engineering. For example, Adams and Le Verrier attempted to explain deviations in the orbit of Uranus by positing an unobserved planet while operating within Newton's theory of gravitation. Similarly, based on theories of elementary particle interactions, astrophysicists hypothesized reaction pathways to explain the relative abundance of elements in stars. Accounts of this sort are more complex than taxonomies or descriptive laws, but their construction is tractable because search is constrained by theoretical knowledge.

Despite these encouraging signs, research on computational scientific discovery still has some important drawbacks. One has been its focus on a single facet of the discovery process – taxonomy formation, law induction, or model construction – while our agents must integrate these abilities. Rare attempts to combine these dif-

ferent elements of science have produced only conceptual demonstrations (e.g., Nordhausen and Langley 1993). Another has been an emphasis on isolated discoveries, rather than the extended research programs our systems must pursue. There have been some efforts along the latter lines (e.g., Kulkarni and Simon 1990; King et al. 2004), but they have been few and far between. Finally, previous work has assumed that observational data has already been collected or that experimental data comes from a fixed laboratory. In contrast, our agents must seek out phenomena in new locations that they select themselves, which raises a number of challenges. We must extend techniques for computational scientific discovery to overcome these limitations before we can embed them in autonomous agents on remote missions.⁵

JOINING EXPLORATION WITH DISCOVERY

I have argued that there is a need for autonomous agents that carry out exploration and discovery in new environments, and also that both research areas are reasonably mature. However, it is not enough to have both functionalities; they must also interact in productive ways. We desire a closed-loop framework, like that as depicted in Figure 1, in which exploration leads to new scientific discoveries that in turn aid exploration. This is similar in spirit to experimental science, but identifying and navigating to new locations involves more than carrying out systematic studies in a laboratory. Let us consider some challenges that the AI research community must address before it can create reliable agents of exploration and discovery.

- *Survival.* Obviously, the autonomous agent must survive while exploring the new environment.⁶ Thus, it should be cautious in moving to new locations and mitigate risk by taking routes it believes are safe. For instance, it might note local fauna follow paths that circumvent rock ledges where dangerous predators might lurk. Similarly, it might rely on generic knowledge to avoid marsh-like regions and stay on apparently firm ground. The agent might also devote its energies to making general observations, rather than spending them on focused studies of specific phenomena. These strategies all rely on knowledge-based reasoning, but they involve very abstract heuristics about responding to possible sources of danger, making them appropriate to a broad range of unfamiliar environments.
- *Environmental protection.* To the extent possible, the agent should avoid damaging or contaminating the environment it explores and studies. Thus, it might favor paths that bypass small or fragile structures, and it might

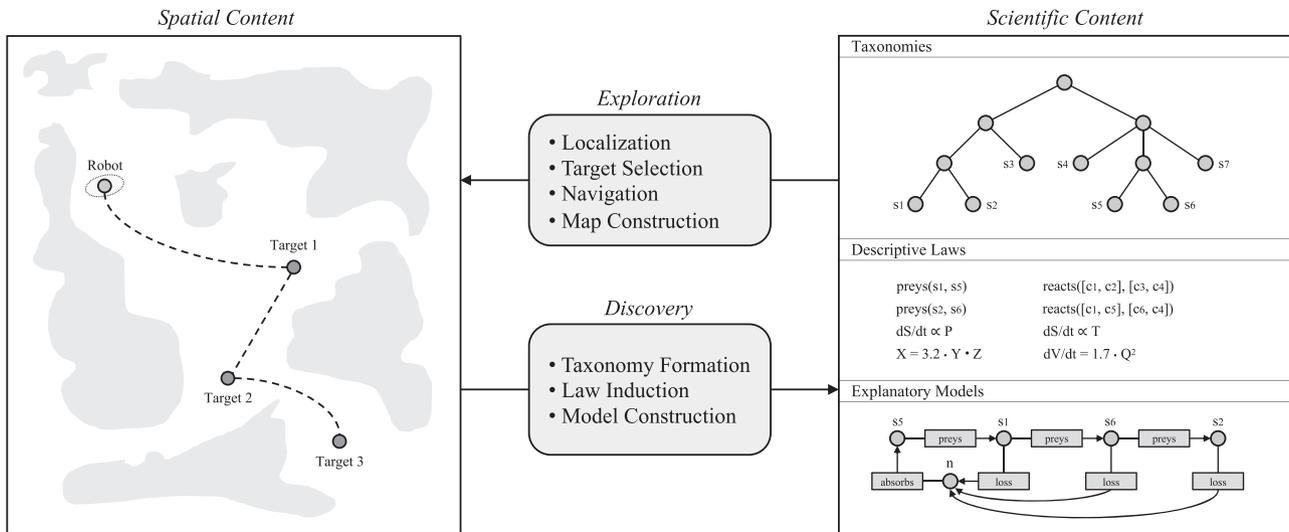


FIGURE 1 The synergistic activities of exploration and discovery. The small boxes itemize the component processes of each activity and the large boxes show the structures they produce. This diagram highlights how the content generated by exploration and discovery inform the complementary process, rather than the influence of spatial and scientific data, which play equally important roles

prefer to observe organisms from a distance, at least until it understands their behavior. There are ethical arguments against interfering with untouched ecosystems or even crystalline formations, but there are also good scientific reasons for minimizing impact on any complex system under study. Any physical agent must interact with the environment it explores, but limiting contacts can extend mission time and avoid contaminating results. Reasoning over generic knowledge can play a role here as well, with an emphasis on survival of the environment rather than that of the agent itself.

- *Scientific interest.* When the autonomous system makes decisions about where to explore, it should take scientific goals into account. This means selecting locations that hold greater potential for novel discoveries and avoiding ones that are well understood. For instance, if the agent observes from afar a seafloor site with distinctive cover, it might favor this choice over ones similar to those already studied. Identifying interesting results has also been an important theme in data mining (Geng and Hamilton 2006), but most work has assumed batch processing of existing data sets. The incremental character of exploratory discovery will let the agent compare new phenomena with current models and detect anomalies that merit further study. This strategy will help guard against confirmation bias and support the cumulative construction of scientific models with ever broader coverage.
- *Focusing attention.* In some environments, the agent must deal with many new entities, which means it must allocate limited attention. This holds even if an agent has an extensive sensor suite that processes its surroundings in parallel, as it must still approach objects to col-

lect samples and run experiments sequentially. Here, the agent might use heuristics similar to those used to select target locations, such as attending to anomalous entities or events that are likely to provide more information. For instance, when collecting rock samples, it might prefer objects that it cannot readily explain with current models, much as surprising results shifted Krebs' attention during his discovery of the urea cycle (Kulkarni and Simon 1990). More generally, the agent must set research goals that are consistent with its high-level mission objectives, then formulate and execute concrete plans to achieve them.

- *Condensing results.* Because the autonomous agent will operate at remote sites with extremely limited communication bandwidth, it cannot simply beam data to home base for processing. Instead, it must filter or summarize the information it collects substantially and convey only the most interesting results (Wagstaff et al. 2005). Fortunately, scientific taxonomies, laws, and models, along with their fit to measurements, offer natural ways to codify the observed regularities in familiar notations that researchers will understand. This includes explanatory accounts, which can often be stated in compact forms that the agent can transmit despite low bandwidth. The most important results will be those that diverge from previous models, suggesting a role for information theory in deciding what to communicate.
- *Increasing effectiveness.* The autonomous system should also take advantage of the knowledge it has discovered to make it a more effective explorer. For instance, the agent might use inferred characteristics of ground cover to follow more efficient paths, draw on regularities in cloud dynamics to make recharging plans, or invoke

newfound chemical reactions to replenish trace minerals needed for its operation. As environmental understanding increases, the agent would get better at the localization, navigation, and maintenance tasks that underlie the activity of extended exploration. This improvement would in turn speed the process of data collection and better support the scientific mission.

- *Collaborating with humans.* I have focused here on autonomous exploration and discovery, but collaboration between intelligent agents and humans also holds great promise, especially in undersea and subterranean settings where two-way communication is occasionally possible. The exploration side has much in common with classic scenarios for human–robot interaction (Goodrich and Schultz 2007), where a remote agent explores and manipulates its environment while a person provides high-level goals and commands. Collaborative discovery can operate in a similar way, where a human scientist selects intriguing phenomena or hypotheses that provide constraints on the AI system’s search (e.g., Bridewell et al. 2006). The ability to consult humans, even rarely, would mitigate the challenges of survival, environmental protection, and attention, making agents for exploration and discovery practical in the near term.

The research community must address these issues before it can deploy autonomous agents to study distant environments. We should also examine multi-agent teams of scientist explorers, which can borrow heavily from earlier results on communication and coordination of distributed systems (Weiss 1999). The ability to share maps, data, and hypotheses will certainly help with some of the challenges outlined above.

Finally, we should devise methods for testing and evaluating the prototype agents that we design and develop. Simulated environments with known landscapes and laws can play a key role, as they could measure the ability to combine exploration with scientific inquiry on extended missions. After this, we can deposit agents in terrestrial environments that are reasonably well understood, such as equatorial rainforests and the Antarctic Peninsula, which will let us test them in more realistic situations. Finally, we can embed our prototypes in uncharted local settings, such as littoral seabeds and underground caverns, that would still allow frequent communication and, if needed, software updates and mechanical repairs.

CLOSING REMARKS

In this paper, I proposed an audacious new challenge for the AI research community: developing autonomous

robotic agents that carry out scientific research in remote and novel environments. I presented motivating scenarios that involved extended missions to distant planetary surfaces and to deep ocean floors, although the idea applies equally well to many other physical settings. Later sections examined two necessary component abilities – *exploration* and *discovery* – in terms of the tasks they address, the constituent processes that support them, and the previous progress in each arena. I argued that both subfields are technologically mature and have contributed to many successful applications, indicating that they could support the vision. However, each paradigm also has some remaining limitations that deserve further attention and effort from relevant researchers and developers.

In addition, I examined some technical hurdles we must overcome to develop and deploy systems that integrate robotic exploration with scientific discovery. These challenges include surviving in unfamiliar and possibly dangerous surroundings, protecting the environment being explored, selecting sites of scientific interest to visit, focusing agent attention on promising phenomena, condensing scientific results for efficient communication, and using these findings to increase agent effectiveness. Although I emphasized autonomous operation, I also noted the potential for human–system collaboration when some remote interaction is possible. Progress on these issues should help the research community to design, implement, and demonstrate the first agents for exploration and discovery in remote and unfamiliar environments. In the longer term, this will lead to wayfaring robotic scientists that follow in the footsteps of Charles Darwin and Alfred Wallace, whose journeys to distant lands resulted in deep insights about life’s evolution.

In closing, I should note that the challenge posed in this paper follows the same spirit that moved the earliest efforts on artificial intelligence. In other words, it identifies a complex cognitive activity that currently only humans exhibit, but that it seems possible, in principle, to replicate. The first stage is to analyze this capability’s components, identify mental structures that can support them, and propose mechanisms that can mimic them. The next step is to consider interactions among the elements and how they can combine to produce the target functionality. This analysis and design stage leads to the implementation of a computational artifact that demonstrates the intellectual behavior on compelling scenarios. Experimental comparisons, organized competitions, or even grand challenges like the one Kitano (2016) has proposed are not appropriate at this stage, as developing even one prototype that exemplifies our vision of a robotic explorer scientist is daunting enough for the AI research community.

ENDNOTES

- ¹ This vision shares some key features with DARPA's SAIL-ON program (Senator 2019) for open-world learning, but the latter differs in its emphasis on responding to changes in environments that are already familiar to agents.
- ² This distinction is not common in AI or other engineering disciplines because researchers and developers have experimental control over their artifacts (Cohen 1995; Kibler and Langley 1988). However, it is widely accepted among historians and philosophers of science, who also study fields that cannot run controlled experiments.
- ³ Some causal models (e.g., Spirtes, Glymour, and Scheines 1993) have an explanatory flavor, but they often relate observed variables and make no reference to domain knowledge, giving them an intermediate status.
- ⁴ Apparently, Joshua Lederberg's enthusiasm about the robotic hunt for Martian life was a major inspiration for DENDRAL. An early report by McCarthy (1964) on the computer-controlled exploration of Mars confirms early AI excitement about this idea.
- ⁵ I have not documented work on taxonomy formation and clustering, which has been an active area since the 1960s, but which has also been largely separate from the efforts on scientific discovery.
- ⁶ This includes monitoring its own health and adapting when problems arise, a topic that I do not have space to discuss here.

ACKNOWLEDGEMENTS

The analysis reported in this paper was supported by the Office of the Under Secretary of Defense (Research and Engineering) and by the Institute for Defense Analyses, which are not responsible for its contents. I thank Daniel Shapiro, Jason Stack, and the reviewers for their constructive comments on earlier drafts.

REFERENCES

- Addis, M., P. C. R. Lane, P. D. Sozou, and F. Gobet, eds. 2019. *Scientific Discovery in the Social Sciences*. Cham, Switzerland: Springer.
- Aha, D. W., M. T. Cox, and H. Muñoz-Avila, eds. 2013. *Goal Reasoning: Papers from the ACS Workshop (Technical Report CS-TR-5029)*. College Park, MD: University of Maryland, Department of Computer Science.
- Anderson, K., E. Bradley, L. Rassbach de Vesine, M. Zreda, and C. Zweck. 2014. "Forensic reasoning and paleoclimatology: Creating a system that works." *Advances in Cognitive Systems* 3: 221–40.
- Arulampalam, M. S., S. Maskell, N. Gordon, and T. Clapp. 2002. "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking." *IEEE Transactions on Signal Processing* 50: 174–88.
- Atanasova, N., L. Todorovski, S. Džeroski, and B. Kompare. 2008. "Application of automated model discovery from data and expert knowledge to a real-world domain: Lake Glumso." *Ecological Modelling* 212: 92–8.
- Beeson, P., J. Modayil, and B. Kuipers. 2010. "Factoring the mapping problem: Mobile robot map-building in the hybrid spatial semantic hierarchy." *International Journal of Robotics Research* 29: 428–59.
- Bresson, G., Z. Alsayed, L. Yu, and S. Glaser. 2017. "Simultaneous localization and mapping: A survey of current trends in autonomous driving." *IEEE Transactions on Intelligent Vehicles* 2: 194–220.
- Bridewell, W., D. Billman, J. N. Sánchez, and P. Langley. 2006. "An interactive environment for the modeling and discovery of scientific knowledge." *International Journal of Human-Computer Studies* 64: 1099–114.
- Bridewell, W., P. Langley, L. Todorovski, and S. Džeroski. 2008. "Inductive process modeling." *Machine Learning* 71: 1–32.
- Bruntona, S. L., J. L. Proctor, and J. N. Kutz. 2016. "Discovering governing equations from data by sparse identification of nonlinear dynamical systems." *Proceedings of the National Academy of Sciences* 113: 3932–7.
- Cohen, P. R. 1995. *Empirical Methods for Artificial Intelligence*. Cambridge, MA: MIT Press.
- Colton, S., A. Bundy, and T. Walsh. 2000. "Automatic identification of mathematical concepts." In *Proceedings of the Seventeenth International Conference on Machine Learning*, 183–90. Stanford, CA: Morgan Kaufmann.
- Durrant-Whyte, H., and T. Bailey. 2006. "Simultaneous localization and mapping: Part I." *IEEE Robotics & Automation Magazine* 13: 99–110.
- Džeroski, S., P. Langley, and L. Todorovski. 2007. "Computational discovery of scientific knowledge." In *Computational Discovery of Scientific Knowledge*, edited by S. Džeroski and L. Todorovski. Berlin: Springer.
- Džeroski, S., and L. Todorovski, eds. 2007. *Computational Discovery of Scientific Knowledge*. Berlin: Springer.
- Ebadi, K., Y. Change, M. Palieri, A. Stephens, A. H. Hatteland, E. Heiden, and A. Thakur. 2020. "LAMP: Large-scale autonomous mapping and positioning for exploration of perceptually-degraded subterranean environments." In *Proceedings of the IEEE International Conference on Robotics and Automation*, 80–6. Paris, France.
- Finn, P., S. Muggleton, D. Page, and A. Srinivasan. 1998. "Pharmacophore discovery using the inductive logic programming system PROGOL." *Machine Learning* 30: 241–70.
- Geng, L., and H. J. Hamilton. 2006. "Interestingness measures for data mining: A survey." *ACM Computing Surveys* 38: 9–es.
- Georgiev, Z., and D. Kazakov. 2015. "Learning ordinary differential equations for macroeconomic modelling." In *Proceedings of the 2015 IEEE Symposium Series on Computational Intelligence*, 905–9. Cape Town, South Africa.
- Geromichalos, D., M. Azkarate, E. Tsardoulis, L. Gerdes, L. Petrou, and C. P. Del Pulgar. 2020. "SLAM for autonomous planetary rovers with global localization." *Journal of Field Robotics* 37: 830–47.
- Goodrich, M., and A. Schultz. 2007. "Human-robot interaction: A survey." *Foundations and Trends in Human-Computer Interaction* 1: 203–75.
- Hidalgo, F., and T. Braunl. 2015. "Review of underwater SLAM techniques." In *Proceedings of the Sixth International Conference on Automation, Robotics and Applications*, 306–11. Queenstown, NZ.
- Ishigami, G., and K. Yoshida, eds. 2021. *Field and Service Robotics: Results of the Twelfth International Conference*. Singapore: Springer Nature.
- Kibler, D., and P. Langley. 1988. "Machine learning as an experimental science." In *Proceedings of the Third European Working Session on Learning*, 81–92. Glasgow: Pittman.

- King, R. D., J. Rowland, S. G. Oliver, M. Young, W. Aubrey, E. Byrne, and M. Liakata, 2009. "The automation of science." *Science* 324: 85–9.
- King, R. D., and A. Srinivasan. 1996. "Prediction of rodent carcinogenicity bioassays from molecular structure using inductive logic programming." *Environmental Health Perspectives* 104(Supplement 5): 1031–40.
- King, R. D., K. E. Whelan, F. M. Jones, P. G. K. Reiser, C. H. Bryant, S. H. Muggleton, D. B. Kell, et al. 2004. "Functional genomic hypothesis generation and experimentation by a robot scientist." *Nature* 427: 247–52.
- Kitano, H. 2016. "Artificial intelligence to win the nobel prize and beyond: Creating the engine for scientific discovery." *AI Magazine* 3(1): 39–49.
- Kocabas, S. 1991. "Conflict resolution as discovery in particle physics." *Machine Learning* 6: 277–309.
- Kostavelis, I., and A. Gasteratos. 2015. "Semantic mapping for mobile robotics tasks: A survey." *Robotics and Autonomous Systems* 66: 86–103.
- Kuhn, T. S. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Kulkarni, D., and H. A. Simon. 1990. "Experimentation in machine discovery." In *Computational Models of Scientific Discovery and Theory Formation*, edited by J. Shrager and P. Langley. San Mateo, CA: Morgan Kaufmann.
- Lane, P. C. R., P. D. Sozou, F. Gobet, and M. Addis. 2016. "Analysing psychological data by evolving computational models." In *Analysis of Large and Complex Data*, edited by A. F. X. Wilhelm and H. A. Kestler. Cham, Switzerland: Springer.
- Langley, P. 1981. "Data-driven discovery of physical laws." *Cognitive Science* 5: 31–54.
- Langley, P. 2000. "The computational support of scientific discovery." *International Journal of Human-Computer Studies* 53: 393–410.
- Langley, P. 2020. "Explanation in cognitive systems." *Advances in Cognitive Systems* 9: 3–12.
- Langley, P., J. Shrager, and K. Saito. 2002. "Computational discovery of communicable scientific knowledge." In *Logical and Computational Aspects of Model-Based Reasoning*, edited by L. Magnani, N. J. Nersessian, and C. Pizzi. Dordrecht: Kluwer Academic.
- Langley, P., H. A. Simon, G. L. Bradshaw, and J. M. Zytkow. 1987. *Scientific Discovery: Computational Explorations of the Creative Processes*. Cambridge, MA: MIT Press.
- Langley, P., and J. M. Zytkow. 1989. "Data-driven approaches to empirical discovery." *Artificial Intelligence* 40: 283–312.
- Li, H., A. V. Savkin, and B. Vucetic. 2020. "Autonomous area exploration and mapping in underground mine environments by unmanned aerial vehicles." *Robotica* 38: 442–56.
- Lindsay, R. K., B. G. Buchanan, E. A. Feigenbaum, and J. Lederberg. 1980. *Applications of Artificial Intelligence for Organic Chemistry: The DENDRAL Project*. New York, NY: McGraw-Hill.
- Mason, S. F. 1962. *A History of the Sciences*. New York: Collier Books.
- McCarthy, J. 1964. "Computer control of a machine for exploring Mars." Stanford Artificial Intelligence Project Memo No. 14. Stanford, CA: Stanford University, Computer Science Division.
- Mitchell, F., D. Sleeman, J. A. Duffy, M. D. Ingram, and R. W. Young. 1997. "Optical basicity of metallurgical slags: A new computer-based system for data visualisation and analysis." *Ironmaking and Steelmaking* 24: 306–20.
- Moravec, H., and A. Elfes. 1985. "High resolution maps from wide angle sonar." In *Proceedings of the IEEE International Conference on Robotics and Automation*, 116–21. St. Louis, MO.
- Nordhausen, B., and P. Langley. 1993. "An integrated framework for empirical discovery." *Machine Learning* 12: 17–47.
- Raissi, M., and G. E. Karniadakis. 2018. "Hidden physics models: Machine learning of nonlinear partial differential equations." *Journal of Computational Physics* 357: 125–41.
- Rosinol, A., M. Abate, Y. Chang, and L. Carlone. 2020. "Kimera: An open-source library for real-time metric-semantic localization and mapping." In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1689–96. Paris, France.
- Senator, T. E. 2019. "Science of AI and learning for open-world novelty (SAIL-ON)." Presented at the Proposers' Day Meeting. Arlington, VA: DARPA. https://www.darpa.mil/attachments/SAIL-ON_Proposers_Day_Distro_A_no_notes.pdf
- Shrager, J., and P. Langley, eds. 1990. *Computational Models of Scientific Discovery and Theory Formation*. San Francisco: Morgan Kaufmann.
- Sleeman, D., V. Corruble, and R. Valdés-Pérez. 2000. "Introduction/editorial: Machine discovery." *International Journal of Human-Computer Studies* 53: 333–4.
- Spirtes, P., C. Glymour, and R. Scheines. 1993. *Causation, Prediction, and Search*. New York: Springer-Verlag.
- Stachniss, C. 2009. *Robotic Mapping and Exploration*. Berlin, Germany: Springer-Verlag.
- Sturges, R. H. 2015. *Practical Field Robotics: A Systems Approach*. Chichester, UK: John Wiley.
- Tanevski, J., L. Todorovski, and S. Džeroski. 2016. "Learning stochastic process-based models of dynamical systems from knowledge and data." *BMC Systems Biology* 10: 30.
- Thrun, S., and M. Montemerlo. 2005. "The GraphSLAM algorithm with applications to large-scale mapping of urban structures." *International Journal on Robotics Research* 25: 403–29.
- Todorovski, L. 2011. "Equation discovery." In *Encyclopedia of Machine Learning*, edited by C. Sammut and G. I. Webb. Boston, MA: Springer.
- Todorovski, L., S. Džeroski, and B. Kompare. 1998. "Modeling and prediction of phytoplankton growth with equation discovery." *Ecological Modelling* 113: 71–81.
- Todorovski, L., S. Džeroski, P. Langley, and C. Potter. 2003. "Using equation discovery to revise an earth ecosystem model of carbon net production." *Ecological Modelling* 170: 141–54.
- Valdés-Pérez, R. E. 1994. "Human/computer interactive elucidation of reaction mechanisms: Application to catalyzed hydrogenolysis of ethane." *Catalysis Letters* 28: 79–87.
- Wagstaff, K. L., R. Castano, S. Chien, A. B. Ivanov, and T. N. Titus. 2005. "An onboard data analysis method to track the seasonal polar caps on Mars." In *Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, 265–72. Munich, Germany.
- Weiss, G., ed. 1999. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. Cambridge, MA: MIT Press.
- Wu, T., and M. Tegmark. 2019. "Toward an artificial intelligence physicist for unsupervised learning." *Physical Review E* 100: 033311.
- Yamauchi, B., and P. Langley. 1997. "Place recognition in dynamic environments." *Journal of Robotic Systems* 14: 107–20.



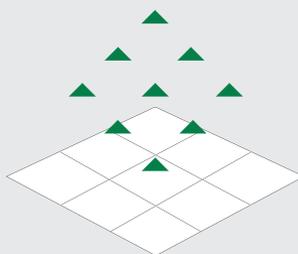
- Yamauchi, B., A. Schultz, and W. Adams. 1998. "Mobile robot exploration and map-building with continuous localization." In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, 3715–20. Leuven, Belgium.
- Zelazny, R. 1975. "Home is the Hangman." In *Analog Science Fiction/Science Fact*. New York: Condé Nast.
- Zhao, W., T. He, A. Y. M. Sani, and T. Yao. 2019. "Review of SLAM techniques for autonomous underwater vehicles." In *Proceedings of the 2019 International Conference on Robotics, Intelligent Control and Artificial Intelligence*, 384–9. Shanghai, China.

AUTHOR BIOGRAPHY

Dr. Pat Langley serves as Director of the Institute for the Study of Learning and Expertise, as a Research Scientist at Stanford University's Center for Design Research, and as an Adjunct Research Staff Member in the Information Technology and Systems Division at the Institute for Defense Analyses. He has contributed to artificial intelligence and cognitive science for more than 40 years, having published over 300

papers and five books on these topics. Dr. Langley developed some of the first computational approaches to scientific knowledge discovery and he was an early champion of experimental studies of machine learning and its application to real-world problems. He is the founding editor of two journals, *Machine Learning* in 1986 and *Advances in Cognitive Systems* in 2012, and he is a Fellow of both AAAI and the Cognitive Science Society. Dr. Langley's current research focuses on architectures for embodied agents, induction of explanatory process models from multivariate time series, and learning complex procedures from written instructions.

How to cite this article: Langley, P. 2021. "Agents of exploration and discovery." *AI Magazine*, 42: 72–82. <https://doi.org/10.1609/aaai.12021>



(Please consult the individual symposium websites for specific submission instructions.)

2022 AAAI Spring Symposium

Monday - Wednesday, March 21-23, 2022

Stanford University, Stanford, California USA

Call for Participation

<https://aaai.org/Symposia/Spring/sss22.php>

Submissions are due November 15, 2021, unless otherwise noted by individual organizers.

The titles of the nine (9) symposia are:

- AI Engineering: Creating Scalable, Human-Centered and Robust AI Systems
- Artificial Intelligence for Synthetic Biology
- Can We Talk? How to Design Multi-Agent Systems in the Absence of Reliable Communications
- Closing the Assessment Loop: Communicating Proficiency and Intent in Human-Robot Teaming
- Designing Artificial Intelligence for Open Worlds
- Ethical Computing: Metrics for Measuring AI's Proficiency and Competency for Ethical Reasoning
- How Fair is Fair? Achieving Wellbeing AI
- Putting AI in the Critical Loop: Assured Trust and Autonomy in Human-Machine Teams