Reasoning and Making Sense of Data in the Absence of Goals

Mary Lou Maher M.MAHER@UNCC.EDU
Software and Information Systems, UNC Charlotte, Charlotte, NC 28202 USA

Kathryn Merrick K.MERRICK@ADFA.EDU.AU
Bing Wang B.WANG@ADFA.EDU.AU
School of Engineering and Information Technology, University of New South Wales, Campbell, ACT, Australia

Abstract

In ill-structured problem domains, it is common to reason about 'problem-finding' or 'hypothesis generation' before goals can be established. Examples of such domains include design, scientific research and mining of 'big data'. Problem-finding and hypothesis generation may also continue throughout the problem-solving process, so identifying goals may be an ongoing process of discovery as well as iterative improvement and refinement. This paper considers the design of cognitive systems that are able to reason in the absence of defined goals. We review a range of approaches that may complement goal-directed reasoning when an artificial system does not or cannot know precisely what it is looking for. We argue that there is a spectrum of approaches that can be used for reasoning or making sense of data in the absence of goals.

1. Introduction

AI approaches to cognitive systems assume that explicit representations of goals, rewards and tasks are integral and provide a focus of attention. Most cognitive systems assume that goals are a starting point for reasoning; that reasoning cannot start without goals; and that reasoning ends when there are no goals. In contrast, it is possible to characterize reasoning so that goals become flexible intermediate structures or implied structures, rather than a predefined and fixed starting point (Maher et al., 2011). While the idea of goals as intermediate structures is similar to agent systems that reason about goals, goal formulation, and goal management (for example, Jaidee et al, 2011), we claim that reasoning in the absence of goals is conceptually different and will lead to different cognitive models. By using concepts such as incentive, novelty, difficulty, complexity, curiosity and surprise, cognitively inspired AI models that mimic human behavior in scenarios such as exploratory design, research and lifelong, self-directed learning are possible. The models can be applied to any domain in which 'problem-finding' needs to occur during problem solving.

The remainder of this section overviews existing approaches to goal-oriented behavior in cognitive systems. The next section examines a number of complementary approaches that may work in conjunction with goal-directed reasoning, including hypothesis generation, motivation, surprise, novelty and curiosity. We classify these approaches along a spectrum that makes progressively weaker assumptions about the definition and presence of goals. We argue that in some cases goal-oriented behavior is an intermediate result of problem-finding, rather than a starting point for problem solving. In such cases goal-oriented behavior can be an emergent property that does not depend on a predefined definition of domain-specific goals.

1.1 Goal-Oriented Behavior in Cognitive Systems

Langley et al. (2008) survey cognitive architectures and layout nine functional capabilities that are required of a cognitive architecture. While each of these are described in terms of functionality without reference to specific architectures or implementations, almost all assume that goals and tasks are inherent in the description of the functions. This can be seen in the way that Langley et al (2008) describe four examples of cognitive architectures:

- Soar: "**All tasks** in Soar are formulated as attempts **to achieve goals**."
- ACT-R: "ACT-R 6 is organized into a set of modules, each of which processes a different type of information. These include sensory modules for visual processing, motor modules for action, an **intentional module for goals**, and a declarative module for long-term declarative knowledge."
- ICARUS: "ICARUS … stores two distinct forms of knowledge. Concepts describe classes of environmental situations in terms of other concepts and percepts, whereas **skills specify how to achieve goals** by decomposing them into ordered subgoals."
- PRODIGY: "On each cycle, PRODIGY uses its control rules to select an operator, binding set, state, or **goal,** to reject them out of hand, or to prefer some over others. In the absence of such control knowledge, the architecture makes choices at random and pursues depth-first means-ends search with backtracking"

There are multiple models for goals – including goal lifecycles and type taxonomies (Braubach et al., 2005) – and processes for solving goals – including machine learning (Nilsson, 1996), planning and rule-based agents (Russell and Norvig, 1995). Braubach et al., (2005) define a lifecycle for goals in which goals transition from new to adopted and finished.

Braubach et al., (2005) also divide goals into a number of types. Approach goals, for example, define states for which an agent should minimize the difference between its current state and the goal state. In contrast, avoidance goals define states for which an agent should maximize the difference between its current state and the goal state. Achievement goals define changes or events that the agent should cause to occur. Maintenance goals define properties that the agent should hold constant. Other types of goals include optimization, test, query and cease goals.

Dignum and Conte (1998) state that truly autonomous, intelligent agents must be capable of creating new goals as well as dropping goals as conditions change. They distinguish between abstract, high-level goals and concrete, achievable goals. They describe goal formation as a process of deriving concrete, achievable goals – such as 'driving at the speed limit' – from high level, abstract goals – such as 'being good'.

Foner and Maes (1994) develop an agent model of unsupervised learning that can self-determine what facts it should pay attention to as a way of modeling focus of attention. Foner and

Maes (1994) distinguish between goal-driven and world-driven focus of attention. In their model, the agent can determine what sensory data to learn from based on strategies that are derived from world-driven goals, such as what has changed recently and what new data is spatially close. These are domain independent strategies that can reduce the number of possible goals an agent can pursue at any given time.

In general, however, there has been less work on how to represent the high-level, abstract goals or world-driven goals that cause new, concrete goals to emerge. The concept of an abstract goal is difficult to formalize because of the difficulty of representing high-level objectives such as "being good" or "being creative". A number of alternative approaches use models of motivation to take the place of abstract learning goals (Merrick and Maher, 2009; Singh et al., 2005; Kaplan and Oudeyer, 2003; Schmidhuber, 1991). Computational models of motivation have also been proposed as an approach to embedding implicit motives in artificial agents to create agents with different preferences for certain kinds of activities (Merrick and Shafi, 2011; Merrick and Shafi 2013). In a different approach, Barnes and Oudeyer (2010) presented a framework for 'maturationally-constrained self-adaptive goal generation' in which an intrinsic motivation module progressively releases constraints on the learning system. This permits the learning system to explore progressively more widely, through the introduction of new goals.

Other work has studied the role of emotion and other cognitive moderators in artificial systems (Mariner and Laird, 2008). Models of emotion act as modifiers to an agent's goal-oriented behavior or provide abstract goals that can be mapped onto concrete goals.

This paper proposes that reasoning can include reasoning before goals are defined, usually based on the current state of the artificial agent and the state of the world. These approaches are consistent with current cognitive systems because they ultimately lead to goal-oriented behavior, but they complement most cognitive systems because they do not assume that goals are predefined. The next section describes models that fall along a spectrum that make progressively weaker assumptions about the definition and presence of goals.

2. Models that Complement Goal-Directed Reasoning

In creative domains such as design and research, the ill-defined nature of tasks suggests a distinction between search and exploration. Maher et al (1996) characterize the difference between search and exploration by the input and output of these processes as illustrated in Figure 1. A typical search process generates a solution as its output with a well-defined problem (or goal) as its input. However, an exploration process derives a problem and the corresponding solution from an ill-defined problem. Maher et al (1996) expand this idea with a co-evolutionary model of reasoning about the problem space and the solutions space in which goals are expressed as requirements in the problem space that are added and adapted in response to the evolutionary search in the solution space.

If we consider the internal state of an agent to include its goals, then the absence of goals remains a valid state. In many autonomous systems the absence of goals implies idle time, but we envisage that a cognitive system can continue to monitor its environment to discover and pursue self-generated goals that extend or improve its knowledge base or skill set, during this so-called idle time. This type of activity will cause changes in the agent's internal and external environment and create a feedback loop that fosters continuous adaptation.

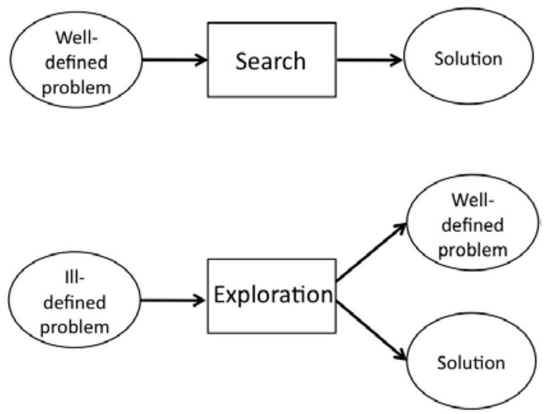

Figure 1. Input and output of search and exploration (Maher et al. 1996).

While ultimately the processes in cognitive systems are organized with the assumption that behavior is goal-directed, we propose that self-directed cognitive systems include the ability to represent, generate, and reason about what Dignum and Conte (1998) call abstract goals. Rather than cast this capability in terms of goals and tasks, however, we identify cognitive models that complement goal-directed reasoning. This is illustrated in Figure 2, where reasoning in the absence of goals can lead to action without an explicit representation of goals or it can lead to the definition of new goals for goal-directed reasoning.

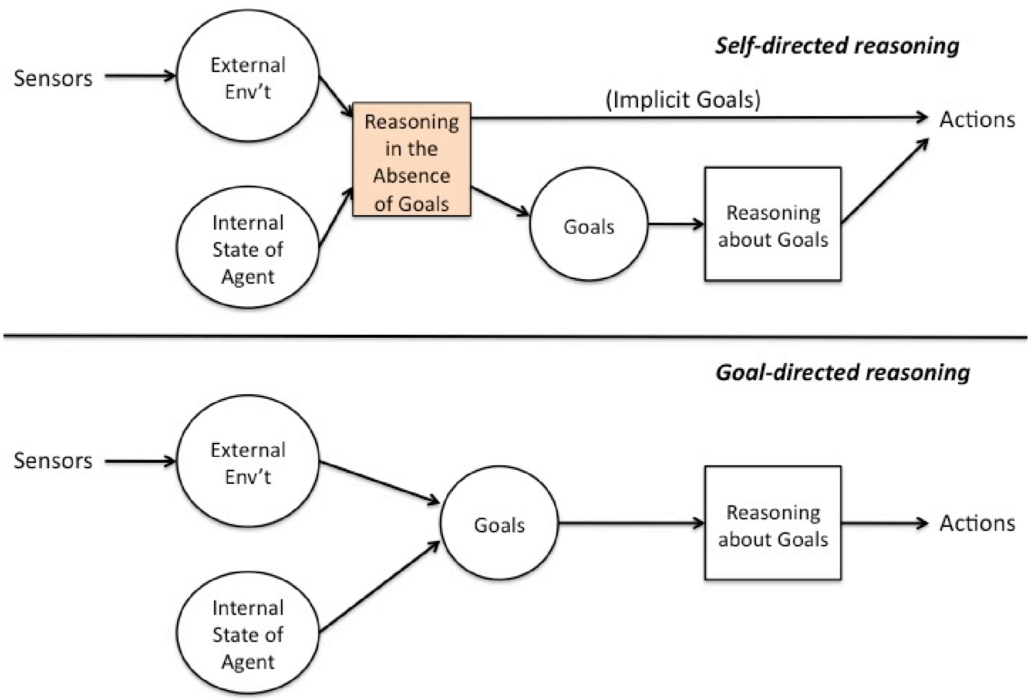

Figure 2: The role of reasoning in the absence of goals in a self-directed cognitive system

In self-directed reasoning, goals are flexible intermediate structures or implied structures, rather than a predefined and fixed starting point for reasoning. Figure 3 shows a spectrum of reasoning starting with the traditional goal-directed reasoning that includes domain specific or state-based goals and models for achieving them, through an intermediate type of reasoning in which goals are implied and may be emergent properties of reasoning, to reasoning without goals in which incentives, for example, provide guidance for reasoning about actions. Goals can be flexible, value-based, intermediate or emergent structures, rather than fixed starting points. This will permit continuous learning and adaptation to unexpected data or events, or changes in needs, beliefs or desires to become an integral concept in cognitive architectures. This will also recast cognitive systems from strictly goal-directed behaviors to include creative and exploratory behaviors.

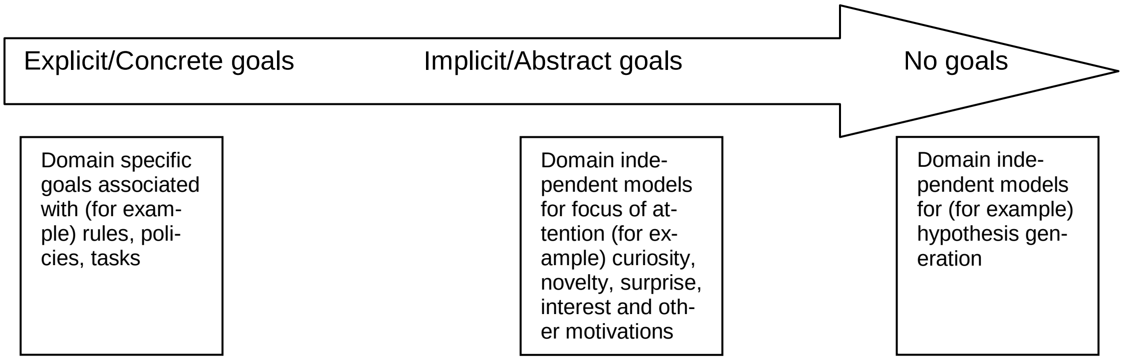

Figure 3. Spectrum of models for reasoning with and without goals.

In this section we begin by looking at models of novelty, curiosity, interest and surprise and the role they play in reasoning where goals are abstract, intermediate or emergent structures. We then consider a set of weaker models based on incentives and motives that represent only a preference for certain types of actions and goals are emergent structures. Finally, we consider models of hypothesis generation that can identify progressively strong patterns and relations in data, when there is no a priori knowledge of the structure of the data.

2.1 Novelty, Interest and Curiosity

Novelty, interest and curiosity fall in a class of models that allow an agent to characterize, act on and learn from changes in the environment. There are many accounts of measuring novelty using computational approaches. Marsland et al. (2000) used Stanley's (1976) model of habituation to implement a real-time novelty detector for mobile robots. Like the Kohonen (1993) Novelty Filter, the real-time novelty detector uses a Self-Organising Map (SOM) as the basis for the detection of novelty. Habituation and recovery extends a novelty filter with the ability to forget.

Models of interest provide a basis for determining if a novel event or state is worth attention. Curiosity is when something of interest can distract the process from its current focus of attention. Saunders and Gero (2001) drew on the work of Berlyne (1960) and Marsland et al (2000) to develop computational models of curiosity and interest based on novelty. They used a real-time novelty detector to implement novelty. Saunders and Gero (2004) model interest using sigmoid functions to represent positive reward for the discovery of novel stimuli and negative reward for the discovery of highly novel stimuli. The resulting computational models of novelty and interest are used in a range of applications including curious agents.

Merrick and Maher (2009) present models of motivated reinforcement learning agents that use novelty and curiosity as models of intrinsic motivation. These agents exhibit a kind of world-driven (rather than goal-driven) behavior. The agents (shown in Figure 4) have an experience trajectory $Y_{(t)}$ that models all states $S_{(t)}$, changes in states (events) $E_{(t)}$, actions that have been encountered/experienced by the agent:

$$Y_{(t)} = S_{(1)}, E_{(1)}, A_{(1)}, S_{(2)}, E_{(2)}, A_{(2)}, \ldots, S_{(t)}, E_{(t)}, A_{(t)}$$

A dynamic motivated reward signal $R_{m(t)}$ is computed as a function of novelty and interest. Their model of interest, based on the experience trajectory, is a modified version of the Saunders and Gero interest function and is based on the Wundt curve shown in Figure 5.

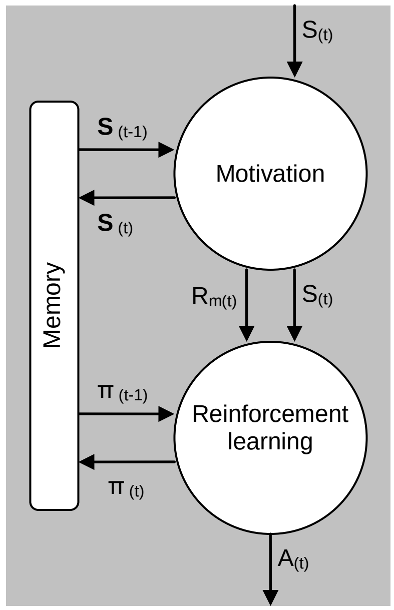

Figure 4 Motivated reinforcement learning agents: goals are implied by a dynamic motivation signal (Merrick and Maher, 2009).

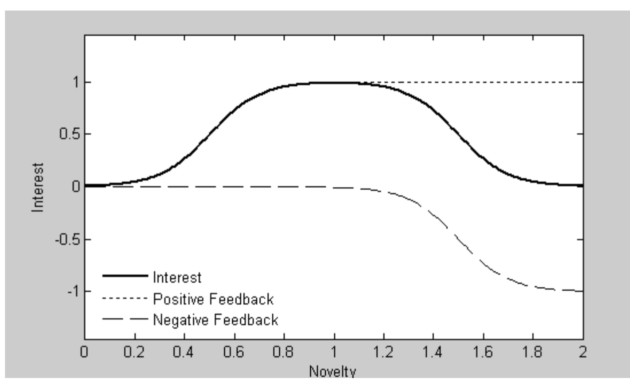

Figure 5 The Wundt curve is the difference between positive and negative feedback functions. It peaks at a moderate degree of novelty (Merrick and Maher, 2009).

This curiosity-based reward signal directs the agent to focus its learning on achieving specific situations at different times, but does not have an explicit representation of tasks or goals. Other motivation functions studied by Merrick and Maher (2009) within this framework include functions for competency and combined competency and curiosity. Experimental studies of curious agents in dynamic environments demonstrated adaptive behaviors through the ability to learn a variety of simple and complex behaviors (see Merrick and Maher, 2009 for experimental results).

2.2 Surprise

Surprise occurs when an unexpected event occurs. While surprise and novelty are similar, something may be novel, but necessarily surprising because it is the next expected change. Horvitz et al (2005) and Itti and Baldi (2004) have developed probabilistic models for finding surprising events in data. Ranasinghe and Shen (2008) have developed a model of surprise for reinforcement learning for developmental robots.

The Horvitz et al (2005) model of surprise is used in traffic forecasting. They generated a set of probabilistic dependencies among a set of random variables, for example linking weather to traffic status. They assume a user model that states that when an event occurs that has less than 2% probability of occurring, it is marked as surprising. Surprising events in the past are collected in a case library of surprises. This provides the data for forecasting surprises based on current traffic conditions. The Itti and Baldi (2004) model of surprise is developed for observing surprising features in image data using a priori and posterior probabilities. Given a user dependent model M of some data, there is a $P(M)$ describing the probability distribution. $P(M|D)$ is the probability distribution after the data is added, using Bayesian probability. Surprise is modeled as the distance d between the prior, $P(M)$, and posterior $P(M|D)$ probabilities.

The Ranasinghe and Shen (2008) model of surprise is used as a reward in a model they call surprise-based learning for developmental robots. In this model, surprise is used to set goals for learning in an unknown environment. The world is modeled as a set of rules, where each rule has the form: *Condition* → *Action* → *Predictions.* A condition is modeled as: *Feature* → *Operator* → *Value.* For example, a condition can be feature1 > value1 where "greater than" is the operator.

A prediction is modeled as *Feature* → *Operator.* For example, a prediction can be "feature1 >" where it is expected that feature1 will increase after the action is performed. The comparison operators provided for surprise analysis include operators to detect the presence (%) or absence (~) of a feature, and the change in the size of a feature (<, <=, =, >=, >). If an observed feature does not match the prediction for the feature, for example, the feature was expected to increase and it decreased, then the system recognizes surprise and sets that state as a reward for learning.

Grace et al (2014) develop a model of surprise based on predictive models using regression analysis or conceptual clustering on product design data. When a new design is encountered that violates our expectations, that design is surprising. This work was further developed to characterize multiple types of expectation and its effect on ways in which we are surprised (Grace and Maher, 2015).

These different approaches to modeling surprise are responses to the needs of the context in which they are developed. The Horvitz et al (2005) model determines that an event in the past is surprising, and then for a collection of surprising events is used to predict future surprising events. In the Itti and Baldi (2004) model, the new data is assimilated into the probability distribution, so something is surprising the first time it is introduced. The Ranasinghe and Shen (2008) model does not use probabilities and instead finds the first unexpected feature based on predictions of the direction in which the values of features will change and sets a reward to learn about that situation. The Grace et al approach is to characterize a cognitive model of expectations in order to recognize when a new design changes our expectations, and is therefore surprising. The role of factors such as curiosity and surprise in information seeking has also been recently considered (Gottlieb et al., 2013).

Surprise serves as a trigger for meta-level reasoning leading to the formation of goals. This concept is explored in the context of design by Grace and Maher (2015).

2.3 Incentives, Motives, and Motivation

In motivational psychology, incentive is defined as a situational characteristic associated with possible satisfaction of a motive (Heckhausen and Heckhausen, 2010). Incentives can be internal or external. Examples of internal incentives that depend on an individual's experiences include the novelty, difficulty or complexity of a situation. Examples of external incentives include money or other kinds of external 'payoff'. Associations between incentive and motivation can be learned, but there are also certain associations between incentives and motivation that have been found to be common across individuals. These include the associations between:

- Task difficulty and achievement motivation
- Risk and power motivation
- Risk and affiliation motivation
- Novelty and curiosity

Suppose we represent a situation encountered by an agent at time t as $S_{(t)}$. Then the incentives associated with a situation can be represented as $I_{(t)} = (i_1, i_2, i_3\ldots)$. Each value i_n represents a different incentive. For example, i_1 may describe the novelty of $S_{(t)}$, i_2 may describe the complexity of $S_{(t)}$, i_3 may describe risk and so on.

Internal incentive values such as novelty, difficulty and complexity can be computed by an agent while it is reasoning about its environment using computational models such as novelty-detectors (Marsland et al., 2000) or achievement based on error calculations on learned policies (Merrick and Maher, 2009). This means that the incentives associated with a situation will change based on the agent's experiences. External incentive values are interpreted from the current state of the environment $S_{(t)}$. These values will change based on changes in the environment. Both types of incentive have the possibility of satisfying the agent's motive.

Implicit motives are innate preferences for certain kinds of incentives. Because different individuals have different implicit motives, they will interpret the same situation incentives differently (Merrick and Shafi, 2013). For example, individuals with strong achievement motivation favor moderate difficulty. Likewise, high curiosity is associated with moderate novelty. Individuals with strong power motivation favor high risk. In contrast, individuals with strong affiliation motivation avoid situations with high risk. We can represent different motives M_1, M_2, $M_3\ldots$ as a function of incentive $M_{m(t)} = M_m(I_{(t)})$. These scalar motivation values $M_{m(\mathrm{t})}$ can be used in isolation, for example as a reward signal in learning, or combined. For example, they can be summed to give a resultant motivational tendency based on a complex motive profile of multiple motives (Merrick and Shafi, 2011).

$$T_{res(t)} = M_{1(t)} + M_{2(t)} + M_{3(t)} + \dots$$

The resultant value can then be used by the agent to identify the most highly motivating situations and act, learn to act or plan to act to achieve those situations. This action, learning or planning may involve formation of explicit concrete goal structures, but this is not strictly necessary.

In summary, models of incentives and motives are able to reason about the synergy of the external environment and the internal state and preferences of the agent to provide a basis for deciding what to do next. They do not represent goal structures although we may recognize emergent goal-directed behavior.

2.4 Autonomous Hypothesis Generation

Advances of computational power, data collection and storage techniques are making large volumes of new data available every day. In some situations, data are collected without a priori supposition or imposition of a specific research goal or hypothesis. Sometimes domain knowledge for this type of problem is also limited. For example, in sensor networks, sensors constantly record data. In these data, expectations about relationships cannot be described in advance. Moreover, the environment may change without a priori knowledge.

While finding patterns in data is achieved with data mining tools and algorithms, there are increasingly situations in which the observational data is collected without specific data mining goals in mind. Big data (Manyika et al., 2011) is a relatively new phenomena that has arisen in cases where very large data sets are accumulated as the result of daily recordings (such as twitter data, phone location data, etc.) for which no specific research purpose was set when the data were recorded. People are interested in analysing the data, however, the questions such an analysis might answer are not initially evident. Hypothesis generation helps to form initial questions that can guide the search for patterns in such accumulated data.

An example is in the field of intelligent systems in which sensors on mobile devices or those embedded in the physical environment record activity data from the environment and can perform pre-defined tasks to adapt to human activities through machine learning techniques. Such tasks can be developed manually when we know the common activities in typical scenarios, e.g., offices or lecture theatres. However, we need new ways for an intelligent environment agent to hypothesise about how to adapt to non-standard scenarios for which knowledge about what the observational data are describing is not available (Merrick et al., 2008).

A similar situation occurs in cyber security in which, due to the constant evolution of hacking activities, previous knowledge about abnormal activities in log data can get outdated. How to use log data to actively acquire updated insights into a system is an interesting research direction for which an agent that can generate new hypotheses from data could be an advantage (Shafi, 2008).

All of these scenarios pose open-ended questions about data. The abstract goal is to make sense of data, but there is an absence of concrete goals to achieve this. We thus argue that our increasing data stores will require new classes of algorithms that tend towards the right hand end of the spectrum shown in Figure 3 and permit reasoning in the absence of goals.

Traditional scientific research (or knowledge discovery) starts with a hypothesis suggesting an interpretation or description of a phenomenon. This hypothesis becomes the foundation for all further inferences and experiments. Its construction is heavily dependent on a researcher's vision and skills, such as observation, domain knowledge, reasoning, imagination and creativity. Once constructed, a hypothesis leads to the design of the experiments and data collection required to test it. Therefore, this type of research is often called hypothesis-testing research (or hypothesis-driven research).

Hypothesis generation (Kell and Oliver 2003; King et al., 2004), which complements conventional hypothesis testing, occurs under circumstances where limited domain knowledge or the size of the data makes it difficult or impossible for a researcher to have well formed expectations or to propose precise hypotheses. There has been research considering hypothesis generation in different fields where hypotheses comprise patterns in the data or machine learning models. Nabel explains this as the methods for hypothesis generation and relates them to the specific environment under study (Nabel, 2009).

In the study carried out by Heintz and Doherty (Heintz and Doherty, 2004), a hypothesis is represented by certain linkage structures, which incrementally grow with additional sensor information being collected. The integration of such structures into agents' functionalities enables the agent to gain awareness of its environment; therefore it provides a basis for other high level goals to emerge. In computer vision, hypotheses often refer to possible matches between two regions of interest (ROI) (Wheler et al. 1995; Chin et al. 2011; Armbruster 2008).

In a computational approach to evaluating the quality of citizen science data, an example of a case in which very large volumes of data are collected, the challenge is determining when an unexpected data item is due to low quality data or if it is the basis for a new hypothesis and therefore is high quality data. This is being explored using concepts from computational creativity and design to search for patterns that are categorized as outliers and harbingers (Maher and Mahzoon, 2015). An outlier is a data item that is unexpected. An outlier is also a harbinger when that data item is the beginning of a new trend in the data, and therefore becomes a pattern that forms a new hypothesis: an indicator to collect more data in that space or to search that space for similar patterns.

Wang et al. (2015) consider the problem of hypothesis generation in continuous data, with the context being situations in which data can be collected about an unknown system. The unknown system is measured by a set of variables. However, limited a priori knowledge is available for characterising the structure and dynamics of system. Wang et. al. define the autonomous hypothesis generation problem in continuous domains in terms of two sub-processes: associative hypothesis generation (AHG) and causal hypothesis generation (CHG). In the case of human knowledge discovery, causation discovery is a progressive process, with the understanding of the causal law behind a system beginning with the observation of associations, which leads to an inquiry into causal relations..

Figure 6 shows the model for hypothesis generation where: The input to the AHG process is X, the set of all data. The output of the AHG process is F, a set of associative hypothesis of the form $X_A \Rightarrow X_B$. The output of the CHG process is G, a causal graph describing relationships among variables in X. Associative hypotheses can potentially reduce the number of variables that need to be examined when forming causal hypotheses and, because their generation procedures exclude irrelevant variables, the CHG can take advantage of the output from AHG to form a causal hypothesis in a reduced variable space. Alternatively, without a priori knowledge about the system, it is possible that there are no specific causal relations between its variables and, if so, it is not necessary to proceed to CHG. In this way, autonomous hypothesis generation can identify progressively stronger patterns in data, when the initial structure and relations in data are unknown (Wang, 2014).

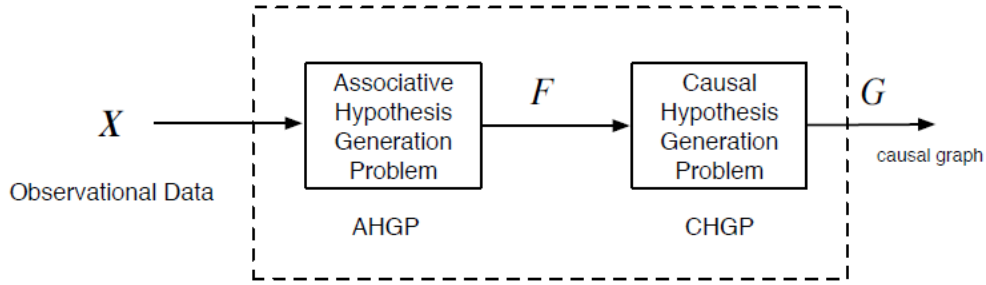

Figure 6 Hypothesis generation identifies progressively stronger hypotheses about data.

3. Summary

This paper has presented several approaches to reasoning in the absence of goals as an alternative to goal reasoning. Introducing computational models of curiosity, interestingness, and motivation lead to alternative models for agents to reason about actions that are not dependent on explicitly represented goals and goal structures. Goals then become intermediate or possibly emergent properties of reasoning. Foner and Maes (1994) distinguish between goal-driven and world-driven attention focus. The models presented in this paper are similar to Foner and Maes (1994) concept of world-driven focus for creating goals or acting when there are no pre-defined explicit goals directing the agent in a learning or problem solving situation. Emergent goal-driven behavior is critical for cognitive systems that can use 'idle time' effectively by continuing to monitor their environment to discover and pursue self-generated goals that extend or improve their knowledge base or skill set.

References

Armbruster W. 2008. Bayesian hypothesis generation and verification. Pattern recognition and image analysis. pp. 269-274.

Baranes, A. and Oudeyer, P.-Y., 2010. Maturationally-constrained competence-based intrinsically motivated learning, IEEE International Conference on Developmenta and Learning, Ann Arbor, Michigan

Berlyne, D., 1960. Conflict, arousal and curiosity. McGraw-Hill, New York.

Braubach, L., Pokahr, A., Moldg, D., Lamersdorf, W., 2005. Goal representation for BDI agents, Second International Workshop on Programming Multiagent Systems: Languages and Tools, p 9-20.

Chin T-J. Yu J. Suter D. 2011. Accelerated hypothesis generation for multi-structure data via preference analysis. IEEE Transaction on pattern analysis and machine intelligence. 99, pp. 1-15.

Dignum, F. and Conte, R., 1998. Intentional agents and goal formation, Intelligent Agents IV: Agent Theories, Architectures and Languages. Springer, pp. 231-243

Fonar L and Maes, P. 1994. Paying attention to what's important: using focus of attention to improve unsupervised learning, Third International Conf on the Simulation of Adaptive Behaviour.

Gottlieb, J., Oudeyer, P-Y., Lopes, M., Baranes, A., 2013. Information Seeking, Curiosity and Attention: Computational and Neural Mechanisms, Trends in Cognitive Science, 17(11), pp. 585-596.

Grace, K., Maher, M. L., Fisher, D. & Brady, K., 2014. Modelling expectation for evaluating surprise in design creativity. In Gero, J.S. and Hanna, S (eds) Proceedings of Design Computing and Cognition 2014, University College London, pp 201-220.

Grace, K. and Maher, M.L. 2015. Specific curiosity as a cause and consequence of transformational creativity, International Conference on Computational Creativity.

Grace and Maher, M.L. 2015. Surprise and reformulation as meta-cognitive processes in creative design, Advances in Cognitive Systems.

Heckhausen, J. and Heckhausen, H., 2010. Motivation and action. Cambridge University Press, New York.

Heintz, F. and Doherty, P., 2004. Managing dynamic object structures using hypothesis generation and validation. In proceedings of the AAAI workshop on Anchoring Symbols to Sensor Data. pp. 1-9.

Horvitz, E., Apacible, J., Sarin, R. and Liao, L. 2005. Prediction, Expectation, and Surprise: Methods, Designs, and Study of a Deployed Traffic Forecasting Service, Proceedings of the Conference on Uncertainty and Artificial Intelligence, AUAI Press.

Itti, L. and Baldi, P. 2004. A Surprising Theory of Attention, IEEE Workshop on Applied Imagery and Pattern Recognition.

Jaidee, U., Muñoz-Avila, H., & Aha, D. W. 2011. Integrated learning for goal-driven autonomy. In Proceedings of the Twenty-Second international joint conference on Artificial Intelligence-Volume Volume Three, AAAI Press, pp. 2450-245I.

Kaplan, F. and Oudeyer, P.-Y., 2003. Motivational principles for visual know-how development. In: C.G. Prince et al. (Editors), Proceedings of the 3rd international workshop on Epigenetic Robotics : Modelling cognitive development in robotic systems, Lund University Cognitive Studies, pp. 73-80.

Kell, D.B. and Oliver, S.J. 2003. Here is the evidence, now what is the hypothesis? The complementary roles of inductive and hypothesis-driven science in the post genomic era. BioEssays, 26:99-105.

Kohonen, T., 1993 Self-organisation and associative memory, Springer, Berlin.

Langley, P., Laird, J.E., Rogers, S. 2008. Cognitive Architectures: Research Issues and Challenges, Cognitive Systems Research.

Maher, M.L., Poon, J. and Boulanger, S. 1996. Formalising Design Exploration as Co-Evolution: A Combined Gene Approach, in J.S.Gero and F. Sudweeks (eds) Advances in Formal Design Methods for CAD, Chapman & Hall, pp 1-28.

Maher, M. L., Merrick, K., Graham, B., 2011. Reasoning in the Absence of Goals, *AAAI Fall Symposium on Advances in Cognitive Systems*, pp 202-209.

Maher, M.L. and Mahzoon, M. J., 2015. *Finding Unexpected Patterns in Citizen Science Contributions Using Innovation Analytics*, Collective Intelligence Conference.

Manyika, J., Chui, M., Brown, B., Bughin, J., Dobbs, R., Roxburgh, C. and Byers, A. 2011. *Big data: The next frontier for innovation, competition and productivity*, McKinsey&Company.

Marinier, R. and Laird, J. E. 2008. Emotion-Driven Reinforcement Learning. CogSci 2008, Washington, D.C

Marsland, S., Nehmzow, U., Shapiro, J, 2000 A real-time novelty detector for a mobile robot, In EUREL European Advanced Robotics Systems Masterclass and Conference.

McClelland, D. C., 1975. Power: the inner experience. Irvington, New York

Merrick, K., and Maher, M. L 2009. Motivated reinforcement learning: curious characters for multiuser games, Berlin, Springer.

Merrick, K. Maher, M.L. and Saunders, R. 2008. Archieving adaptable behaviour in intelligent room using curious supervised learning agents. In Proceedings of CAADRiA 2008 Beyond Computer Aided Design, pages 185-192.

Merrick, K., Shafi, K., 2013. A Game Theoretic Framework for Incentive-Based Models of Intrinsic Motivation in Artificial Systems , *Frontiers in Cognitive Science, Special Issue on Intrinsic Motivations and Open-Ended Development in Animals, Humans and Robots.* Baldassarre, G., Barto, A., Mirolli, M., Redgrave, P., Ryan, R., Stafford, T (Eds). Volume 4, 30th October, 2013

Merrick, K. and Shafi, K., 2011. Achievement, affiliation and power: motive profiles for artificial agents. Adaptive Behavior, 9(1): 40-62.

Nabel G. 2009. The coordinate of truth. Science, 326, pp. 53-54.

Ranasinghe N. and Shen,W-M, 2008. Surprise-Based Learning for Developmental Robotics. In Proc. 2008 ECSIS Symposium on Learning and Adaptive Behaviors for Robotic Systems, Edinburgh, Scotland.

Russell, S. J. and Norvig, P. 1995. Artificial intelligence: a modern approach, Prentice Hall, Englewood Cliffs, NJ

Saunders, R., 2001. Curious design agents and artificial creativity. PhD Thesis, University of Sydney, Sydney

Saunders, R. and Gero, J.S., 2004. Curious agents and situated design evaluations. Artificial Intelligence for Engineering Design, Analysis and Manufacturing, 18(2): 153-161

Schmidhuber, J. 1991 A possibility for implementing curiosity and boredom in model-building neural controllers. In The International Conference on Simulation of Adaptive Behaviour: From Animals to Animats pp 222-227

Shafi. K. 2008. An online and adaptive signature-based approach for intrusion detection using learning classifer system. PhD thesis, University of New South Wales Canberra Campus.

Singh, S., Barto, A.G. and Chentanez, N., 2005. Intrinsically motivated reinforcement learning. In: L. Saul, Y. Weiss and L. Bottou (Editors), Advances in Neural Information Processing Systems (NIPS). The MIT Press, Vancouver, pp. 1281-1288

Stanley, J. C., 1976. Computer simulation of a model of habituation. Nature 261:146-148

Wang, B., Merrick, K., Abbass, H., 2015, Autonomous Hypothesis Generation as an Environment Learning Mechanism for Agent Design. In Proceedings of the 2015 Conference on Artificial Life and Computational Intelligence, pp 210-225.

Wang, B, 2014, Autonomous Hypothesis Generation for Knowledge Discovery in Continuous Domains, PhD Thesis, University of New South Wales, Canberra.

Wheler M. and Ikeuchi K. 1995. Sensor Modeling, Probabilistic hypothesis generation, and robust localization of object recognition. IEEE Transaction on pattern analysis and machine intelligence. 17(3), pp. 252-265, 1995.