Water resource planning and management using motivated machine learning

JANUSZ STARZYK
School of Electrical Engineering and Computer Science, Ohio University, Athens, Ohio 45701, USA
starzykj@ohio.edu

Abstract Water resources planning and management require problem resolution and optimized use of resources. Since many objectives in water management are conflicting, it is hard to devise one optimum strategy. A simulation tool capable of optimized multi-objective analysis to satisfy a multiplicity of goals is needed to support water decision making. This paper suggests an integrated modelling framework to assist with time consuming and difficult tasks of decision making by water management practitioners and to harmonize economic uses of water resources. Motivated machine learning, presented in this paper, supports intelligent decision-making processes in dynamically changing environments and could be used to consider alternative water management policies. Motivated learning systems learn to properly control the environment with competing goals. They provide a natural support for multi-objective decision making in an active search for balance between conflicting situations and adverse environmental conditions. A case study of optimized machine learning water management decisions is presented.

Key words multi-objective analysis; water management; dynamic environments; motivated learning; competing goals; goal creation

INTRODUCTION

Planning and management of water resources is based on the shared interests of all the stakeholders involved and has a significant effect on their economies and quality of life. Despite the increasing importance of modelling in water resources planning and management, no single tool or methodology provides a satisfactory solution. There is uncertainty in how to use data, how to develop infrastructure, which objectives must be satisfied, how to share resources and how to plan and coordinate joint efforts. The available models are mostly limited to regional-level strategies, while the challenges are transdisciplinary and involve expertise from many sciences and engineering fields. Water allocation between conflicting uses and among competitive users is becoming an important issue. Intense competition arises between countries and between users, even within the same country during a lean season. Uncertainty of future changes in climate and water uses, the complexity of water-related effects on the environment, health and development, economy and policy making are critical to reaching acceptable solutions.

Water resources systems have many mutually dependent factors and are at the heart of economic policy making between the nations that share water resources. Managers and planners involved in developing these policies must identify and evaluate alternative designs. There is a growing demand for integrating mutual dependencies of water sciences and policy making. Therefore, a need for legitimate scientific data and integrated optimization tools to estimate values of decision variables that best support specified objectives has been emphasized repeatedly.

Water is needed to satisfy the basic human needs such as hygiene, drinking, cooking, farming and recreation. The quality of water must be protected by proper legislatures, monitoring, and enforcement to prevent water-borne diseases, prevent water contamination, and protect ecosystems. Ordinary people suffer if everyone involved tries to extract as much water as one can get and degrades the water resource. Water resource management involves monitoring and management of water quality and quantity, and making plans and predictions related to water usage and supply, water storage and purification, contamination control and law enforcement. This raises a number of questions, such as what policies and legal framework can promote sustainable use of water, or how to protect water resources from overuse and contamination?

Understanding of natural processes as well as their social and economic impact on society, services, or resources provided by rivers, lakes, reservoirs and wetlands is important for successful water management and policy making. These processes involve many physical and biological systems interdependencies whose understanding is at the core of successful decision making and planning. It has become increasingly evident that the water problem has become too complex, interconnected and large to be handled by any one institution or group of professionals, irrespective of their competence, authority and good intentions.

Decision makers are ready to consider the uncertainty that results from model imperfections and their impact on the predicted outcomes of the decisions made. Promoting agreeable-to-all decision-making procedures and incorporation of uncertainty into model parameters are critical to economic development and security in developing societies. Current management regimes for determination of what is good for people can fail and result in poor water quality and related society ills. Monitoring the impact of water management policies and execution of effective monitoring plans are vital to assess how well the policies are meeting people's expectations. This may require setting proper guidelines for measuring performance indicators, such that the decision makers can be informed about the results of their actions and may adapt their policies to satisfy changing requirements.

Water resources affect a variety of economic, environmental, and ecological issues and may result in social tension and affect the quality of life. They also serve other useful roles, such as drought and flood control, hydro-energy production, transport, recreation and waste disposal. Local community roles and initiatives in water management practices are a key to improving socially critical aspects of insufficient water supply and improving water quality. Participation of local communities in water management results in better organization and more effective utilization of water than the services provided by governmental agencies. Cooperative management to facilitate water use seems inescapable, since such cooperation is critical for sustainable use of water.

Various performance criteria can be used to provide measures of how well the water-related problems are addressed by the system. Due to a variety of goals some of these criteria may be in conflict with others and various trade-offs must be considered. It is a challenge to evolve strategies for equitable and sustainable water use by creating frameworks for rapid policy changes. The relations between the society and the policy makers have become polarized around water resources, and it is necessary to create assessment and modelling tools to improve policy making and facilitate interaction.

Computer simulation, modelling and decision support tools provide only a partial guidance to the decision-making process by synthesizing and optimizing results of quantitative analyses based on the provided data and modelling objectives. Various models have been developed to assess and predict water supply and water quality and their impact on economic and social development. Advanced computerized models include optimization methods, fuzzy sets, genetic programming, data mining and artificial neural networks. Machine simulation of model uncertainty involves statistical methods, probability estimation, stochastic modelling and sensitivity analysis.

For instance, a decision support system for water management discussed by Muleta & Nicklow (2005) combined evolutionary algorithms with a watershed simulation model to arrive at recommended land use through the solution of a multi-objective optimization problem. The model used a soil and water assessment tool for evaluation of the objective function and used the strength of the evolutionary algorithms to handle multiple objectives. This model can be applied on various watershed scales. Msiza *et al.* (2008) used computational intelligence techniques for modelling and prediction of short-term and long-term water demands. The paper used a growing power of computational intelligence to model the dynamically changing environment and decision-making process.

This work improves on the existing computerized methods, proposing a machine learning approach, known as motivated learning, to support the decision-making process. It addresses critical issues of water resource planning and management, such as model and data uncertainty, dynamic changes in the environment, development of policy making process, or handling of competing and often conflicting goals.

MOTIVATED MACHINE LEARNING MODELS

Water resource management tools include simulation, optimization and multi-objective analysis. The question is how to design a simulation tool to support water decision-making satisfying a multiplicity of goals including multi-objective decisions. Computerized models were used for many years to support water-related decision making and water resource management. However, model development is limited to the expertise of those who propose them, and as a result models often overly simplify dynamics of economic, social and environmental interactions and lead to inappropriate policy making and management decisions. A different approach is proposed in this paper, in which models are replaced with real, dynamically changing environments with all of their complex intricacies and societal dependencies. This idea has been successfully applied to the development of autonomous robots which interact with dynamically changing environments and learn proper interactions without building environment models, but rather using the environment the way it is.

The main objective of this paper is to suggest an integrated modelling framework that may assist with the time-consuming and difficult tasks of decision making by water management practitioners, and harmonization of economic uses of water resources. An integrated and effective machine learning platform may help to build effective partnerships between modellers and practitioners in the development and application of water management models and observe them in handling simulated crisis situations. Motivated machine learning that provides seamless support for intelligent decision-making processes in dynamically changing environments could be applied to consider alternative water management policies. It may be able to incorporate socio-cultural, political, economic and institutional elements that influence decision making, addressing non-dominated solutions.

This framework uses a goal creation approach in embodied intelligence (EI) that motivates the machine to develop into a useful research tool through active interaction with the real environment. It integrates modelling with planning, decision making, policy implementation and evaluation, using dynamic feedback from the field to modify models and the decision-making process. The method adapts to changes in the environment conditions, and resistance to policy implementation, and human factors, showing robustness under uncertain parameters, imperfect data, and imperfect models.

In the motivated learning (ML) systems, different types of data received from the environment are associated and represented to build knowledge and the environmental model. This representation is validated through active interaction with the environment. Learning in such systems is incremental, with constant prediction of the input associations based on the emerging models and only new information is registered in the system memory. Knowledge is not entered into such systems, but rather is a result of their successful use in a given environment. ML provides natural support for multi-objective decision making, focusing on the most pressing issues in an active search for balance between conflicting situations and adverse environmental conditions.

Motivated intelligent systems adapt to unpredictable and dynamic situations in the environment by learning, which gives them a high degree of autonomy, making them a perfect choice to support human decision making (Pfeifer & Bongard, 2007). This approach described by Starzyk (2008) uses emerging, self-organizing, goal creation (GC) systems that motivate embodied intelligence to learn how to efficiently interact with the environment. The motivated learning mechanism was designed to provide motivations to the learning machine that combine its externally driven goals with internal goals that emerged from the developmental process and are controlled internally by the machine. Motivated learning first learns dependencies between objects in the environment and the externally set objectives (controlled by the external rewards), and subsequently, uses these observations to set internal goals. Most of the time the machine responds to specific goals, trying to find solutions to the problem set, so it explores the environment with a specific objective. Motivated learning uses negative reward systems as its major reinforcement. Negative signals (also known as pain signals) are received from the environment and need to be minimized (synonym of reward). If the negative signals increase instead of being reduced, this corresponds to a negative reward and the machine learns not to perform actions that resulted in such an increase.

A ML machine is in a continuous process of building new motivations and responding to established ones. Competing signals that represent abstract pains and attention-switching direct the machine to choose a goal to act on and to follow this goal. These signals vary as the machine acts and the environment around it changes. The mechanism to build motivations and choose goals, triggers learning of intentional representations and establishes associations between sensory observations and motor actions. The machine responds to the observed environment changes and to its own internally generated pain signals to choose the proper action. This response is as much a result of top-down deliberation and prediction of what will be the result of its action, as of bottom-up perceptions, experiences and past history of interactions with the environments.

It is suggested to further extend the ML approach and to apply it to practical and theoretical aspects of water management in changing environments, where the existing methods fail or work with difficulty. For instance, using classical machine learning to predict the future for physical processes works only under the assumption that the same processes will repeat. However, if a process changes beyond certain limits, the prediction will not be correct. The expectation is that ML systems may overcome this difficulty and that such systems can be trained to advise humans.

CREATION OF ABSTRACT MOTIVATIONS AND ABSTRACT GOALS

An abstract pain signal is created once the machine is unable to perform the action that resulted in the reduction of the lower-level pain. For instance, if a machine needed a certain resource to alleviate its primitive pain, and the resource is not available or is hard to find, this creates an abstract pain signal. This abstract pain motivates the machine to explore how to obtain the missing resource. An abstract pain centre is not stimulated from a physical pain sensor; it only symbolizes the pain of not having the resource that the machine needs to prevent the primitive pain.

Suppose that an agent receives several “primitive” pain signals that indicate that he is “dirty”, “thirsty”, or has discovered “drought” (see Fig. 1). Depending on which signal dominates, the agent tries to lower this pain. For instance, if the agent is thirsty, he can learn that drinking water will lower his primitive pain. However, when there is no more water, he cannot alleviate this pain. Therefore, the agent develops an abstract pain related to the lack of water. Once created, this pain centre will compete with other pains for attention, independently of the original primitive pain that was responsible for its creation. Thus, an abstract pain leads to a new learned motivation that may direct an agent to perform certain actions independently of the primitive pain. For example, an agent may not be thirsty, and yet, if there is no water, he will look for it whenever this abstract pain “lack of water” dominates.

Motivated by this new abstract pain, the agent needs to learn how to overcome it. It may find out that it can draw water from the well. Thus, it learns a new concept (the well) and is able to recognize the well as something related to its needs (specifically lack of water). It learns a new useful action “drawing the water from the well”, and it associates this action with the means to remove his abstract pain of not having water (its abstract goal). It also expects that, after performing this action (drawing water from the well), it will get water. This expectation will be useful for future action planning. At the same time, another higher-level pain (and motivation to remove this pain) develops related to the availability of drawing water from the well. Thus, if the well it was exploiting dries out, or the agent no longer can access this well, it may need to learn to overcome this pain, for instance by digging a new well. An alternative to this solution could be to maintain the proper level of the groundwater and this higher-level need may lead to the need for resource utilization planning by regional managers or policy makers. This process can be illustrated using Fig. 1.

The network of motivations can be expanded both vertically (towards a higher abstraction level) as well as horizontally (on the same abstraction level). For instance, rather than drawing water from the well, an agent may learn that it is easier to buy water. It develops an alternative way to accomplish this goal. This will lead to an understanding of a new concept (money) and new abilities (buying water). Related to this will be another abstract pain of not having enough money, and related means to get rid of this pain (like developing ecotourism, rising taxes, or providing services like digging a well). While some motivations may point to new higher-order motivations (like developing infrastructure to attract tourists), others may point to motivations previously developed both on a higher level (digging a well) or lower level (using a water reservoir for water recreation facilities).

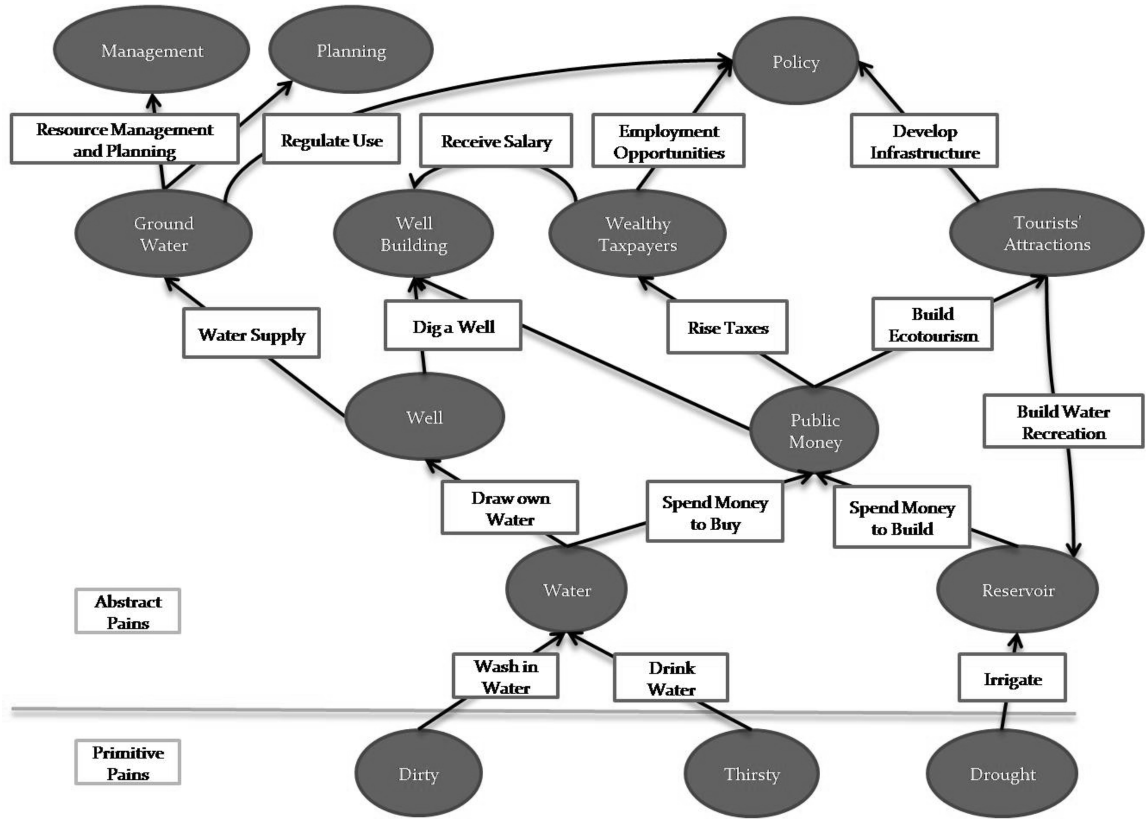

Fig. 1 Creation of abstract pain signals.

Notice, that in the presented scheme, some goals may provide a circular path. For example, the need for the water reservoir was an abstract pain developed through the action of irrigating fields to remove the primitive pain signal related to drought. An abstract motivation resulting from the need for a water reservoir could be to earn money (to pay for the reservoir), while building tourists attractions may be a way to earn money, and finally building a new reservoir may be motivated by the need to develop tourist attractions. Thus, a learning network must be able to detect and avoid using such circular solutions. In the proposed motivated learning scheme, this is accomplished by blocking the circular goals through inhibitory, unsuccessful action neurons.

The machine is motivated by competing pain signals to act and to discover new ways of improving its interaction with the environment. By doing so, the machine not only learns complex relationships between concepts, resources and actions; it also learns limitations of its own embodiment, and effective ways of using and developing its motor abilities. The machine learns to associate its motivations with goals that lead to deliberate actions. It learns the meaning of concepts and objects, and relations among objects, learns to perform new actions and to expect results of actions. This builds up complex motivations and higher-level goals as well as the means of their implementation. Based on competing pain signals, the machine chooses which actions to execute to satisfy its goals and manages the goal priorities at any given time.

The motivated learning may provide a useful tool to support decision making, planning and management of water resources, as it develops a natural way of balancing various competing needs. Discovery of new ways to accomplish specific objectives gives the machine freedom to decide how to approach a given problem. Unlike other machine learning methods it does not follow a prescribed algorithm to optimize its decision, it does not require the environmental model but it is capable of building one, and it dynamically adjusts its actions to changes in the environment and its own perception of needs.

The most advanced machine learning method used so far in developmental robotics and autonomous systems control uses principles of reinforcement learning (RL) (Bakker &

Schmidhuber, 2004). In the next section we compare motivated learning with reinforcement learning under identical environment conditions.

COMPARISON BETWEEN ML AND RL APPROACHES

We compared the performance of learning agents based on motivated learning and reinforcement learning principles in an identical environment that has many dependencies between its resources, as illustrated in Fig. 1. In our experimental set-up the agent has to operate in an environment where not all resources can be used from the very beginning of simulation. Instead, gradually more and more resources are available to the agent during the interaction with the environment. At the beginning of simulation the agent is able to learn only the basic dependencies between resources. It can also choose from a small set of actions. Additionally, it has a given a period of time before the environment increased in complexity introducing other resources and making new actions available.

The basic concept is that when the agent uses some resource, the amount of the resource decreases and in order to replenish it the agent has to choose and perform a proper action which may use another "higher-level" resource. Moreover, there are only small amounts of various resources, and while the agent consumes them, its pain signals P_p increase. The agent should learn which actions to perform in order to replenish the resource which is needed at this very moment. The following function describes the probability of finding resources in this experimental set-up:

$$f_{ci}\left(k_{ci}\right)=e^{-\frac{k_c}{\tau_c}} \tag{1}$$

where τ_c is the scaling factor that describes a resource declining rate, and k_c the number of times a resource was used.

The environment's state changes as a result of actions performed by the agent. As the agent uses up the resources from the environment, the resources are harder to find unless the agent learns how to restore them. The best strategy for an agent in such a situation is to learn about the environment when it is still simple.

Our aim in this study is to use this environment with gradually increasing complexity to examine the effectiveness of motivated learning (ML) and compare it with reinforcement learning (RL). The results are shown in Figs 2 and 3. Figure 2 shows the moving average of the pain signal value, P_P, while Fig. 3 shows the ratio of the pain signal values in RL and ML as a function of the number of iterations.

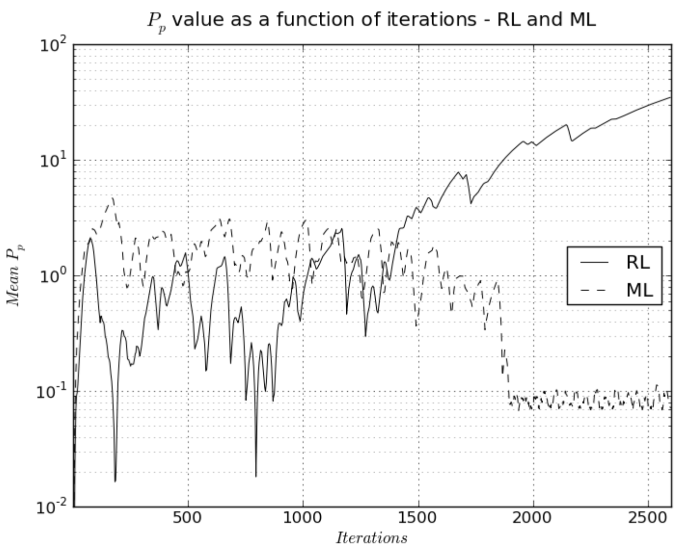

Fig. 2 Moving average of P_P value as a function of number of iterations.

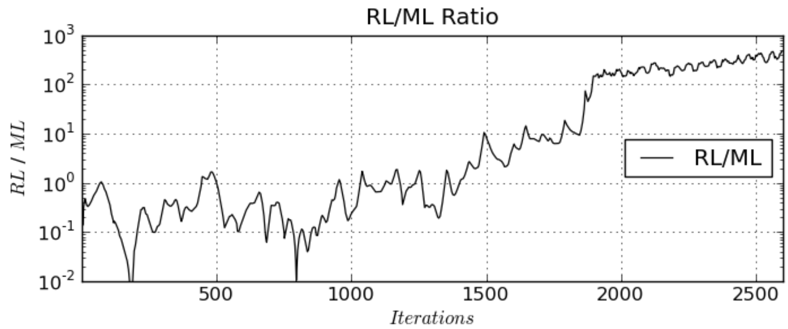

Fig. 3 RL/ML P_P ratio in experiment with 14 competing goals.

As we can see, the ML-based agent was able to converge to a stable solution with low pain signals, P_p, while the RL-based agent cannot do that in this kind of environment. While initially RL was able to maintain the average pain level (even below that of ML), it gradually performed worse as the environment conditions deteriorated. As expected, learning is more effective with gradual use of new resources and skills in an increasingly complex environment. This follows intuition: initial environment simplicity should result in quicker learning because the learner's efforts are not diffused by different possibilities.

CONCLUSIONS

This paper presents the motivated learning approach and its potential use in water resource planning and management. The proposed ML method shows how system development stimulates learning of new concepts and at the same time benefits from this learning. ML yields a machine learning system that may support monitoring and optimization of the system performance, while choosing an action to address the most pressing problems from many that may compete for the manager's attention. It is expected that this natural learning will lead to more accurate models for water-related policies and actions, and through active interaction with human expertise will use the provided input data identifying factors that most contribute to water supply, use, contamination or policy-making decisions.

A case study of machine learning water management decisions is presented in this paper to demonstrate the application of EI in facilitating humans with modelling and water-related decision-making process.

REFERENCES

Bakker, B. & Schmidhuber, J. (2004) Hierarchical reinforcement learning with subpolicies specializing for learned subgoals. In: *Proc. 2nd IASTED Int. Conf. on Neural Networks and Computational Intelligence* (ed. by M. H. Hamza), 125–130. Grindelwald, Switzerland.

Msiza, I. S., Nelwamondo, F. V. & Marwala, T. (2008) Water demand prediction using artificial neural networks and support vector regression. *J. Computers* **3**(11).

Muleta, M. K. & Nicklow, J. W. (2005) Decision support for watershed management using evolutionary algorithms. *J. Water Resour. Plan. Manage*. **131**(1), 35–44.

Pfeifer, R. & Bongard, J. C. (2007) *How the Body Shapes the Way We Think: A New View of Intelligence*. The MIT Press (Bradford Books), Cambridge, USA.

Starzyk, J. A. (2008) Motivation in embodied intelligence. In: *Frontiers in Robotics, Automation and Control* (ed. by A. Zemliak), 83–110. I-Tech Education and Publishing, USA.