

dynamically samples the input data online so as to excise long sequences of repetitive data that would otherwise prevent shorter critical sequences from being properly learned.

The method of vector quantization employed by our governor architecture is Resource Allocating Vector Quantization (RAVQ). The RAVQ dynamically generates model vectors that represent important subsequences in the input stream. We use the RAVQ classification to identify critical periods in the sensory data by training the neural network controller on sensory inputs when the classification produced by the RAVQ changes. This classification change corresponds to a critical period in the robot's sensory environment. Long sequences of repetitive data, primarily responsible for catastrophic forgetting, do not cause the RAVQ to change classification, and thus are dynamically excised from the training set.

The experiment consisted of training a neural network robot controller on a wall following task with both governed and ungoverned inputs. Networks in both cases were trained for 20,000 time steps and evaluated for 1,000 time steps. Networks were scored based on the total sum of their distance from the wall over the evaluation period. The teacher was also evaluated to serve as a baseline for comparison. The number of successful trials using the governed network exceeds that of the ungoverned network. The mean score across all successful trials is also significantly lower.

For this simple task, using vector quantization to avoid classification proved to be far superior to training on unbalanced raw data. The ability to employ back-propagation learning algorithms for robot control tasks enables the construction of hierarchical developmental systems. By autonomously extracting key information from the input data, vector quantization may serve as an appropriate solution in many more circumstances where balanced, critical event-based data needs to be extracted from large datasets.

MATHEMATICAL MODELS GOVERNING DYNAMIC

USER-DRIVEN RESOURCE ALLOCATION*

STUDENT POSTER ABSTRACT

Michael Massimi, massimi2@tcnj.edu

Eric Tarn, tarn2@tcnj.edu

The College of New Jersey

Advisor: Dr. Ursula Wolz

* Copyright © 2004 by the Consortium for Computing Sciences in Colleges. Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the CCSC copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Consortium for Computing Sciences in Colleges. To copy otherwise, or to republish, requires a fee and/or specific permission.

In many complex computer-driven tasks, select information is hidden from users, either on purpose or inadvertently. When this is the case, users may become frustrated because they are not alerted to the causes of delays or other malfunctions. In turn, overall efficiency of a system can decrease while users wait for the system to time out or otherwise resolve the delay. One way to minimize the ill effects of resource conflicts is to involve users in the process of determining resource allocation by providing a forum for self-governance and problem resolution.

We posit a voting system as a method of conflict resolution. We use voting to solve a particular task - that of allocating fixed computing resources (e.g. bandwidth, CPU time, disk space, etc.) in small, ad hoc networks. Users are involved in the decision-making process. When users join the network, they are given a neutral priority rating for resource usage. They are also given a mechanism for calling a "referendum" on any user's rating, including their own. Involving the users in such governance will not only enable people to come to a communally agreeable solution to the conflict, but will also increase the overall satisfaction of the users within the system.

A linear voting system has the potential to allow users to dominate the resource and cause undesirable situations. We developed a voting formula that permits us to maintain the integrity of the system, while offering users the chance to become contributing members. We imposed boundaries on the maximum and minimum ratings that users may attain, and the formulae determine the distribution of the ratings within those boundaries. By imposing boundaries, we ensured that all users have a high enough priority to successfully access the resource; thus, no user is subject to starvation.

The formulae used to express ratings (see below) include mathematical references to the greater community. The two values that comprise a user's new rating (T_{new}) are the impact (I) and the change (C). The impact is the ratio of the ratings of the source (S) and target (T) users. This captures the relative credibility of the involved users to determine the extent to which the change will be applied. The change places all user ratings under a bell curve variant. The median (M) of all other users is compared to the target user's rating (T), and then placed into an integrated curve defined by the maximum rating change (R). We are testing these formulae under best, worst, and average case conditions.

The proliferation of shared resources demands systems such as ours to help users access them efficiently and politely. Our future work will include thorough examination of the aggregate wait time for all users involved in a set of transactions to determine if allowing people to resolve their conflicts a priori will result in a more communally efficient solution.