

COMPUTATIONAL RESOURCE MANAGEMENT IN SUPERVISED LEARNING SYSTEMS

Louis A. Tamburino
Wright-Patterson AFB

Mateen M. Rizki
Wright State University

Michael Zmuda
Wright State University

ABSTRACT

The allocation of computational resources is an important design issue in systems involving the exploration of large search spaces. In this paper, we explore a novel strategy for adaptively adjusting a resource allocation control parameter. This adaptive strategy modifies the control parameter using minimum a priori information. We use the technique in a specific supervised learning system and evaluate it by comparing the values generated by the control strategy to optimal control parameter values obtained from extensive measurements.

SUPERVISED LEARNING PROCESSES

Supervised learning systems, which operate in large search spaces, require efficient resource allocation strategies in order to maximize the performance achievable with fixed amounts of total resources. Strategy is driven by performance measures which provide feedback information during the search process.

The search process explored in this paper consists of a sequence of independent experiments, each of which is characterized by a certain number of operations. The objective of the search is to find a detector that will separate target images from non-target images.

Associated with each detector generated in the search procedure is a performance measure P which is normalized such that $Y = P/P_{\max}$ ($0 < Y \leq 1.0$). The search process is limited to a total resource allocation, T . The search process consists of sequences of independent experiments in which different detectors are generated. Let $D(k)$ denote the k -th detector generated with an allocation of $N(k)$ operations and $Y(k)$ the corresponding terminal performance. The larger the assignments for $N(k)$, the fewer the number of independent experiments because one is limited to T operations for all experiments.

Statistically, the higher the value of $N(k)$ used to generate $D(k)$, the higher the expected value for $Y(k)$. However, by assigning a value larger than necessary, one reduces the number of experimental trials; this in turn reduces the expected optimal performance with T . On the other hand, if the values of $N(k)$ are lower than what is required to generate an acceptable level of performance, T operations are wasted on a large number of experimental trials, each having an inadequate assignment of operations $N(k)$.

ADAPTIVE CONTROL STRATEGY

We present a resource allocation strategy which adjusts the values of $N(k)$ to converge to an optimal value. The optimal ranges for the control parameters in our experiments have been measured by systematically assigning a very large number of operations to different values of $N(k)$ and recording the optimal performances. These measurements are discussed below and are used to evaluate our control algorithm. The optimal range varies for different training sets. The algorithm must therefore quickly estimate the optimal range using performance information from the initial experiments. The use of statistical analysis is precluded because the initial feedback information is relatively sparse. The approach described below is a feasible alternative: a dynamic averaging process which provides the estimates for optimal settings for $N(k)$.

Our strategy is based on a dynamic envelope model in which the value of $N(k)$ is adjusted by feedback performance information during the course of a search. During the generation of each $D(k)$, we monitor the performance as a function of the number of operations expended, $N \leq N(k)$. These observations can be visualized as trajectories on a Y - N graph where performance is plotted as a monotonically increasing function of operations expended up to the limit $N(k)$. The basic idea is that of a self-adjusting envelope which contains or covers all observed trajectories. The

envelope then serves as a reference for estimating the optimal setting for $N(k+1)$ to be used in the next experiment.

Optimal settings maximize results in terms of the largest performance measure attained during a search using T operations. If the values of $N(k)$ are quickly made to converge and remain within an optimal range of values, then the probability of achieving maximum or near maximum performance for some detector is increased.

The dynamic envelope used in the algorithm consists of two straight line segments: a horizontal segment S_h and a diagonal segment S_d . The equations are given by

$$S_h: Y = H$$

$$S_d: Y = H + m \cdot (X - X_v),$$

where H , X_v , and m are parameters that define S_1 , and

$$X = \ln N / \ln T, \quad N = T^{**} X \text{ (number of operations).}$$

The two line segments intersect at the envelope's vertex V which is located at point (X_v, H) in figure 1. The slope of the diagonal segment, specified by m , is taken to be a constant.

The region of optimal performance for the control parameter X_r is located to the right of the envelope vertex, V . In our initial experiments, we adjusted the control parameter as follows:

$$X_r(k) = X_v(k) + r(H),$$

where the length of the horizontal segment r varies with H :

$$r(H) = (1 - H)/(2 + H).$$




Figure 1. Control Envelope.

Recall that $N(k) = T^{**} X_r(k)$.

Note that the exponential parameter X_r varies between 0 and 1. Typically, there is an upper limit $X_{max} < 1.0$ which guarantees a minimum number of experimental detectors $D(k)$ and there is a minimum value of $X_{min} > 0$ necessary to register a non-zero performance value. If $X_r = 1$, then all resources will be expended on one experiment, $D(1)$, which may not necessarily produce an acceptable detector.

The horizontal segment is supported by the maximum observed performance. If a new maximum value for H is established, the vertex moves upwards. When an experimental trajectory penetrates the envelope to the right of the vertex, the vertex is moved to the right. When a trajectory penetrates the diagonal segment, the vertex is moved to the left. These opposing events are not balanced. The probability of crossing the diagonal segment is greater than that of crossing the horizontal segment. The net effect is to drive and hold the envelope too far to the left. In order to establish a balance, a constant bias or incremental movement to the right is used to pull the vertex gradually to the right. This bias movement is eventually stopped by trajectories that can penetrate the diagonal segment of the envelope more easily as it shifts to the right. In this way, the envelope vertex adjusts itself based on the evolving performance and establishes an equilibrium position such that $X_r(k)$ converges to a desirable range of values.

EXPERIMENTAL RESULTS

A simple pattern recognition problem serves as the test-bed for the envelope control strategy. In this problem, detectors are hit-or-miss templates that discriminate among uppercase letters of the English alphabet. An individual template is formed through an evolutionary process that refines the template by adding points. As points are added, the cost of the template is increased.

Results are presented for training experiments using the letters E, F, X, and Y. The total number of operations, T , for each experiment is four million resource units or word operations (WOpts). The angle of the diagonal segment of the envelope is 60° . The initial setting of resource units per detector is two thousand ($X(1) = 0.5$). A population of detectors is generated within the resource limitations and whenever the performance measure of a detector under construction penetrates the control envelope, X_v is adjusted. T is divided into twenty-five measurement intervals. At the end of each interval, the position of X_r is recorded. The position of X_r at these sample times defines a control history. The experiment is repeated twenty-four times to verify the results (see figure 2a-d).

A control history consists of a period of initial adjustment followed by a persistent span of banded oscillations. During the adjustment period, the value of X_r fluctuates widely from its starting point of 0.5. The length of this interval varies for different target letters. For the letters E and F, the initial intervals are clearly visible; they extend from sample time 1 to sample time 5 (see figures 2a, 2b). In contrast, the adjustment periods for the letters X and Y are not clear-cut, but they appear to range from sample time 1 to sample time 10 (see figures 2c, 2d).

After the adjustment period, a typical control history enters an interval of restricted oscillation. In figure 2a, the banded oscillation extends from 5,000 word operations to 10,000 word operations. The width of this band suggests that an expenditure of approximately 7,500 word

operations per detector provides the necessary and sufficient quantity of resources to generate a detector that achieves maximum performance. The location and range of the band of oscillations varies for different target classes. For the letter F, the band stretches from 7,500 to 18,000 word operations, while the bands for the letters X and Y oscillate between 10,000 and 35,000 word operations (see figures 2c-d).

The envelope-based control strategy is robust. The behavior of individual control histories is consistent among the twenty-four independent experiments for each target letter. These families of histories terminate their adjustment phase at approximately the same time and oscillate with a similar amplitude.




Figure 2a.




Figure 2b.




Figure 2c.




Figure 2d.

Figures 2a-d. 60° Envelope Control History. Each figure represents 24 independent experiments. Each of these trials used

4,000,000 word operations. The location of X_r is shown at discrete resource intervals.

There are three parameters associated with the envelope control strategy: the slope m of the diagonal S_d , the distance from the vertex X_v to the selected resource allocation level X_r , and the level of decay. To investigate the sensitivity of control to variations in the angle of the diagonal segment, the experiments are repeated using a 45° segment (see figures 3a-d). The qualitative features of the trajectories observed using a 45° diagonal are the same as those obtained using the 60° diagonal. There is an initial adjustment interval followed by a stage of banded oscillation.

There are some differences, however. The rate of convergence to the banded oscillation is slower for control histories generated using a 45° segment. The

histories generated for the letter F are the most striking example of this behavior. With a 45° diagonal, the initial adjustment period extends to sample time nine (figure 3b) in contrast to sample time 5 for the 60° diagonal (figure 2b). Also, the amplitude of the variations after convergence is greater for the trajectories established by the 45° segment. The best example of this is evident in the histories generated for the letter Y (see figures 2d and 3d). For the 60° diagonal, the band of oscillation is 10,000 to 35,000 word operations while the range of oscillation for the 45° diagonal stretches from 5,000 to 45,000 word operations. Further experiments are under way to test additional slopes as well as a strategy that employs a self-adjusting slope.




Figure 3a.




Figure 3b.




Figure 3c.




Figure 3d.

Figures 3a-d. 45° Envelope Control History. Each figure represents 24 independent experiments. Each of these trials used

4,000,000 word operations. The location of X_r is shown at discrete resource intervals.

EVALUATION

The control histories generated for each letter suggest a range of acceptable settings for word operations per detector. To verify the accuracy of these ranges, a maximum performance curve is computed for each letter (figures 4a-d). Each curve summarizes the results of 312 experiments distributed across 13 resource allocation levels. A single experiment consists of generating a population of detectors and recording the maximum performance achieved by the best detector after 4,000,000 word operations are spent. The performance

curves log the average, minimum, and maximum performances attained by the best detector in 24 independent experiments at the 13 different levels of resource allocation. The cost to construct a single maximum performance curve is 1.258 billion word operations. In contrast, the control algorithm is able to estimate the region of maximum performance while using only a fraction of the allotted four million word operations.

The performance curves clearly reveal a desirable range of resource allocation levels. For the letter E, levels of resource allocation of approximately 1,000 to 10,000




Figure 4a.




Figure 4b.




Figure 4c.




Figure 4d.

Figures 4a-d. Maximum Performances. Each plotted point represents the results of 24 experiments using a fixed number of resources per detector. The detector with the best performance in

each experiment is used to obtain the minimum, maximum, and average for each level of resource allocation.

WOpt per detector produce the highest minimum, average, and maximum performance (figure 4a). Allocation levels below 5,000 do not consistently produce top performers as evidenced by the lower averages and minimums, while settings beyond 20,000 WOpt show a dramatic drop in the average and minimum performances. The envelope control strategy using the 60° diagonal segment suggests that an allocation level of 5,000 to 10,000 WOpt per detector is sufficient (see figure 1a).

The performance curve for the letter F suggests that the acceptable level extends from 5,000 to 40,000 resource allocation units (figure 4b). Again, this is consistent with the control strategy recommendation of 7,500 to 20,000 WOpt (figure 2b). The performance curves indicate that levels between 10,000 and 100,000 are favorable for generating X detectors (figure 4c) while the control strategy predicts a range from 10,000 to 35,000 (figure 2c). Finally, the performance curve for the letter Y indicate that a range from 5,000 to 60,000 word operations per detector is acceptable (figure 4d), and the control strategy specifies a range from 10,000 to 35,000 (figure 2d).

DISCUSSION

The values recommended by the control strategy correspond closely to the values obtained through extensive experiments. The strategy predicts parameter setting at the low end of the desired range but the predicted value is always beyond the allocation level required to produce a detector with maximum performance. The strategy favored parameter values that are just below the average of the maximum performance curve.

The parameter settings produced by the algorithm using the 45° diagonal segment (see figures 3a-d) are consistent with the optimum values displayed in the maximum performance curves (figures 4a-d). The bands of oscillation using the 45° diagonal are slightly larger than those observed using the 60° diagonal and overlap the 60° bands. Therefore, the control algorithm predicts similar settings for both 45° and 60° diagonals.

Our results demonstrate that the envelope control strategy selects good parameter settings with minimum a priori information. The strategy requires the specification or adjustment of three parameters: the distance from the envelope vertex to X_r , the envelope decay rate, and the angle of the diagonal segment. Experiments to explore the sensitivity of the control strategy to changes in these parameters are under way. Preliminary results using the 45° and 60° diagonal segments suggest that the strategy is relatively insensitive to slope. Additional work using self-adjusting angles is planned.

ACKNOWLEDGEMENT

This work is supported by the Air Force Office of Scientific Research, Directorate of Electronic and Material Sciences under task 2305R5 and contracts F33615-87-C-1416 and F49620-88-C-0053.