

This program is due 27 November 2007 in class.

The goal of this assignment is to extend your knowledge of machine learning. In a previous assignment, you used hill climbing to learn the appropriate weights for a neural net backgammon player. In this assignment, you will use a special form of reinforcement learning called $TD(\lambda)$.

1 TD-Gammon Revisited

As mentioned in your programming assignment 3, Tesauro created a master backgammon player he called TD-Gammon. Here are some additional details derived from

<http://www.univ.kiev.ua/yawd/books/AI/11/node2.html>
<http://www.cs.ualberta.ca/sutton/td-backprop-pseudo-code.text>

Tesauro first created a neural network player he called *Neurogammon*. Neurogammon had special features encoded into the neural network for representing backgammon playing knowledge and was trained against hand-picked examples. It learned to play reasonably well and was categorized as an “intermediate” player.

Tesauro’s next experiment was a backgammon player that learned by playing itself with zero-knowledge of the game, except for required technical information such as legal moves, win-states, etc.. Tesauro applied a learning method described as a “gradient-descent form of the $TD(\lambda)$ algorithm” that we will describe later. This player, called TD-Gammon 0.0 trained by playing itself. After 300,000 rounds, it could play as well as all contemporary computer programs.

Later modifications reincorporated the specialized backgammon features into the neural network (TD-Gammon 1.0) and multi-ply searches (TD-Gammon 2.0 and 3.0). The most recent version plays at the level of the best human champions in the world.

In this assignment, we will be re-creating TD-Gammon 0.0 for your intellectual stimulation and pleasure. Additionally, we will be examining both the theory of the method and the practical minutia that results from the application thereof. It should be noted that training takes time, so **START EARLY**.

2 The TD algorithm

Before we discuss the TD algorithm, we should talk briefly about the supervised learning technique known as *backpropagation*. Backpropagation works for neural networks that are “feed forward”, meaning there is no feedback or loops. Provided with an expected output, backpropagation takes the error and applies it starting at the outputs and propagating backwards with the amount of correction dependent on the existing weight.

But for backpropagation to work in a game like backgammon, after every move the network would require the “correct” decision and use that knowledge to train the network. As that information is not available, Tesauro used temporal difference methods. TD methods are somewhat like backpropagation *over time* to assign credit or blame of some reward to a previous state.

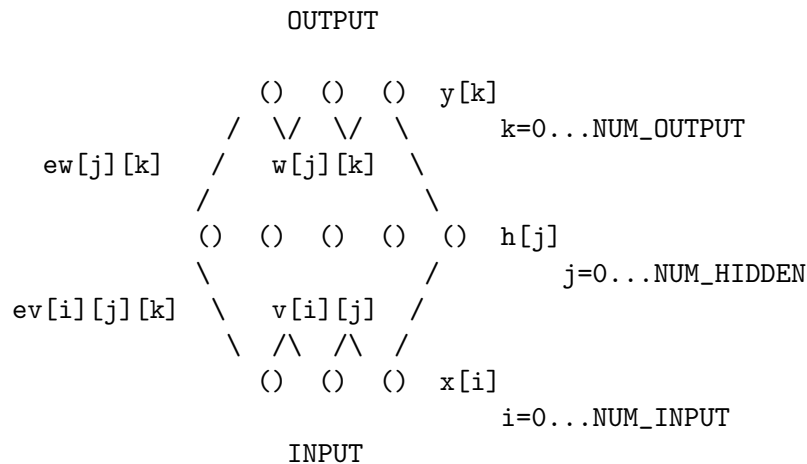
More specifically, when the network wins or loses a game, were all moves equally responsible for that win or loss? Obviously not. So, a temporal difference method is used to identify and reinforce the credit or blame with respect to the outcome.

The actual mechanics work like this. You start with a random network. This network is used to play *both sides* of a backgammon game. At each point, the network evaluates the next moves to pick the best move. Obviously, when the weights are random, the moves will also be random. However, even though the evaluation means nothing at the start, you still treat the next evaluation *as if it were the correct one*. So, for instance, if it is white’s move, and the evaluation says that white’s odds of winning are 0.75, and then on black’s move, the evaluation function recalculates the new odds of white winning to be 0.80, the algorithm treats the 0.80 to be the correct value, *even though the new number is also an estimate*, and trains the network accordingly. This component behaves exactly like backpropagation.

What the temporal difference algorithm adds is to maintain information over time about these trainings. While the information does decay eventually, it does flow forward in time influencing future training. When one player or the other actually wins, and the odds of winning are now concrete (i.e., 0.0 or 1.0), a more correct error is computed and applied to training *with influence from the previous trainings!*

Even though the data is flowing forward in time, the net result is that the impact of winning or losing is backpropagated to the previous moves played.

The structure of the neural network we are using is shown below:



Mathematically, at each move, the weights are adjusted according to the following formula:

$$w_{jk}^{t+1} = w_{jk}^t + \beta \cdot \Omega_k \cdot e_{jk}^t$$

$$v_{ij}^{t+1} = v_{ij}^t + \alpha \sum_k \Omega_k \cdot e_{ijk}^t$$

where

- w_{jk} is the weight from hidden unit j to output unit k
- v_{ij} is the weight from input unit i to hidden unit j
- k is an output unit
- $\Omega_k = y_k^{t+1} - y_k^t$ is the difference of the output of node k between time $t + 1$ and t
- α, β are learning rates
- e_{ijk}, e_{jk} are the *eligibility traces* at time t

Eligibility traces represent the information carried forward from one training to another. They are computed using the following formula:

$$e_{jk}^{t+1} = \lambda e_{jk}^t + (y_k^t \times (1 - y_k^t) \times h_j^t)$$

and

$$e_{ijk}^{t+1} = \lambda e_{ijk}^t + (y_k^t \times (1 - y_k^t) \times w_{jk}^t \times h_j^t \times (1 - h_j^t) \times x_i^t)$$

where

- λ is the decay parameter
- y_k^t is the value of output node k at time t
- h_j^t is the value of hidden node j at time t
- x_i^t is the value of input node i at time t

It can take some time to understand these equations, so **START EARLY**.

3 Implementation Details

Despite a plethora of published research about or based on TD gammon, there is actually very little written about how to recreate Tesauro's results. Even Tesauro's paper itself is sketchy on details.

Here are the known configuration parameters:

- $\alpha = 0.1$
- $\lambda = 0.7$
- **Hidden Unit Count = 40**

There are at least two unknowns.

1. **Did Tesauro wipe out his eligibility trace data at the end of each complete game?**⁴
Theoretically, an eligibility trace should only apply to a sequence of moves within the same game. However, experimentation on the part of the TA suggests that it works better to not reset the data each round.
2. **How did Tesauro create his evaluation function?** Tesauro had four outputs; two represented the odds of white or black winning, and two represented the odds of white or black getting a *gammon*. So, if the output suggested that the odds of white winning was 0.8 and the odds of black winning was 0.8, how should it evaluate the board?

Additionally, three additional tweaks are often used in reinforcement learning that Tesauro does not mention.

1. **Alpha Splitting.** The first tweak is splitting α into two parameters (α and β) and applying one learning on the weights from input to hidden units, and the other for learning on the weights from hidden units to output nodes.
2. **Alpha Decay.** The second tweak is to use an α decay. The idea is that at the beginning of training, α should be large leading to large changes. But, as time goes on, α should decay to prevent the accidental loss of good playing knowledge.
3. **Player History.** It is possible for a good player to be “lost” by accident during training. Backgammon involves some chance, and poor behavior might occasionally appear good. One option is to save the player’s neural network and restore it if it’s performance worsens.

One of your tasks in this assignment is to experiment with three of the following parameters: α , λ , hidden unit count, output count (and representation), evaluation function (only with multiple outputs), eligibility trace lifetime, alpha splitting, alpha decay, or player history. You should try a number of combinations of values for your three parameters to explore the impact.

It can take a long time to run all of these experiments, so **START EARLY**.

Here are a few other hints.

- **Neural Network Setup.** For this assignment, you need to follow Tesauro’s parameters for the *inputs* to the network. That means you need 198 units.
 - For each point on the backgammon board, four inputs represent how many pieces of one color are on that point. If there is 0-3 pieces, then 0-3 inputs should be enabled (e.g., 0 pieces = 000, 1 = 100, 2 = 110, 3 = 111) and if there are n pieces, with $n > 4$, then the first three units are set to 1, and the fourth unit is set to $(n - 3)/2$. This accounts for 192 units
 - Two additional units encode the number of white and black pieces on the bar where the value of the input is $n/2$
 - Two more units encode the number of white and black pieces successfully removed from the board. The value of each input is $n/15$.
 - Finally, two units encode whose turn it is to play. For white, 01 and for black 10.

- **Shared Network.** When you're playing rounds using the code framework, both play⁵ers are sharing the *same neural network and training parameters*. In the TA code, our `BackPropPlayer` has a method `createLearningPartner` that returns another `BackPropPlayer` that shares the network.
- **Player Evaluation.** Your player will be evaluated by its performance against your expecti-minimax player from programming assignment #2. Every 1000 rounds, you should test your player and report the number of wins out of 100. For fun, you might also test against a random player and your hill-climbing neural network player. If you feel that your expecti-minimax player is too weak (or is easily beaten by your td-gammon player), contact the TA for a stronger opponent, or attempt to integrate with gnubg.
- **Rounds.** You should run your experiments for at least 30,000 rounds. If you are really ambitious, try running them for 300,000 rounds. You should know that 300,000 rounds will probably require a day or two, so if you want to do that, **start early**.
- **Non-Learning Mode.** You need to be able to turn learning off for when your player is playing evaluation rounds.
- **Unexpected Results.** The TA code has, in various configurations, experienced improvement in the first 50,000 rounds, and then slow worsening as the rounds approach 200,000. This may be a bug in our code, or it may be expected. If you aren't sure if your player is learning, contact the TA for help.

4 Deliverables

You are expected to submit the following:

1. The code for a working player that learns from self play. Your code must conform to the following specification (assuming that `Main.java` is your central class).


```
java Main -N [rounds] -O [neural network output] -I [neural network input]
```

 Note that we expect your player to be able to save and load neural network files. We are not distributing a framework for this assignment. You should be able to use the framework you used for programming assignment #2, but it will take some time to get everything working, so **START EARLY**.
2. The neural network output of your "best" player.
3. A written submission. Your reports should include your significant implementation choices and some thoughtful discussion, as well as any specific data requested in the assignment. Specifically, you must include:
 - (a) Graphs showing the learning of your player. The x axis is thousands of rounds, the y axis is number of wins out of 100 against your expecti-minimax player. There should be at least six graphs, one for at least two settings of your three experimental parameters. You'll need more graphs if you do combinations. You should automate the generation

of your graph. Your code can spit out the results of its play every thousand rounds into a file for processing by gnuplot.

- (b) A discussion of why you think your parameters had the effect that they did.
- (c) A discussion of why did or why did not this learning method work better than hill climbing (you may need to test your hill climbing player against this player, or against expecti-minimax for comparison).
- (d) A discussion of why you think that TD learning works for backgammon and how well you think it would work for other games and/or AI related problems.

As converting this mathematical description into code is difficult, we provide you with a pseudo code implementation of the procedure. However, it will probably take some time to debug and get working correctly, so **START EARLY**.

```

/* MOST OF THIS IS ACTUAL JAVA CODE, BUT BEWARE      */
/* IT is STILL PSEUDO-CODE-ish... YOU WILL HAVE      */
/* TO FILL IN THE BLANKS (make real variables, etc.) */

public double gradient(HiddenUnit u) {
    return u.getValue() * (1.0 - u.getValue());
}

/* Ew and Ev must be set up somewhere to the proper size and set to 0 */
public void backprop(double[] in, double[] out, double[] expected) {
    /* compute eligibility traces */
    for (int j = 0; j < net.hidden[0].length; j++)
        for (int k = 0; k < out.length; k++) {
            /* ew[j][k] = (lambda * ew[j][k]) + (gradient(k)*hidden_j) */
            Ew[j][k] =
                (LAMBDA * Ew[j][k]) +
                (gradient(net.hidden[1][k]) * net.hidden[0][j].getValue());
            for (int i = 0; i < in.length; i++)
                /* ev[i][j][k] =
                    (lambda * ev[i][j][k]) +
                    (gradient(k)+w[j][k]+gradient(j)+input_i)*/
                    Ev[i][j][k] = ( ( LAMBDA * Ev[i][j][k] ) +
                        ( gradient(net.hidden[1][k]) *
                            net.hidden[1][k].weights[j] *
                            gradient(net.hidden[0][j])      *
                            input[i]
                        )
                    );
        }

    double error[] = new double[out.length];

    for (int k =0; k < out.length; k++)
        error[k] = expected[k] - out[k];

    for (int j = 0; j < net.hidden[0].length; j++)
        for (int k = 0; k < out.length; k++) {
            /* weight from j to k, shown with learning param of BETA */
            net.hidden[1][k].weights[j] += BETA * error[k] * Ew[j][k];
        }
}

```

```

        for (int i = 0; i < in.length; i ++) {
            net.hidden[0][j].weights[i] += ALPHA * error[k] * Ev[i][j][k];
        }
    }
}

public Move move(Backgammon backgammon) {
    Move bestMove = null;
    double expectedUtility = -1.0;
    double nextoutput [] = null;
    for (Move m: backgammon.getMoves()) {
        /*
         * output is an array of 1-4 depending on configuration
         *
         * evaluate the next board for the OTHER PLAYER; you'll
         * have to adjust your utility accordingly...
         *
         * getCurrentBoard is the board AFTER the move
         */
        double output[] = net.getValue(m.getCurrentBoard(), otherPlayer);
        double utility = computeUtility(output);
        if (utility > expectedUtility) {
            bestMove = m;
            expectedUtility = utility;
            nextoutput = output;
        }
    }
    if (learningMode) {
        /* you'll need to pass the input values to backprop... */
        double currentinput[] = boardToVector(m.getOriginalBoard());

        /* get original board is the board BEFORE the move... i.e., right now */
        /*
         * OH OH!!!! WARNING WARNING WARNING, DANGER WILL ROBINSON!!!
         * you must call the neural net's get value on the original input
         * BEFORE you call backprop. The HiddenUnit class has a "getValue" that
         * is cached; it is not computed every time. If you don't call this
         * before backprop, then the nodes will have the values set from
         * the last call to the whole network's getValue() which may or may not
         * be the one you want
         */
        double currentoutput[] = net.getValue(m.getOriginalBoard(), thisPlayer);

        /*

```



```

    * notice, you don't have to monkey with nextoutput even though it is
    * for the other player... this is because it's output is player independent.
    * if one output, it equals odds of white winning, if two outputs than
    * one is odds of white winning, other is odds of black winning
    */
    backprop(currentinput[], currentoutput[], nextoutput[]);
}
return bestMove;
}

public void lost() {
    if (learningMode) {
        /* assuming two output with odds for white and black */
        double actual[] = new double[2];
        if (myPlayer == Board.WHITE){
            actual[Board.WHITE]=0.0;
            actual[Board.BLACK]=1.0;
        }
        else{
            actual[Board.WHITE]=1.0;
            actual[Board.BLACK]=0.0;
        }

        /* you'll have to save currentBoard from the move method or something */
        double in[] = boardToVector(currentBoard);
        double out[] = net.getValue(in);

        backprop(in, out, actual);
    }
}

/* figure out boardToVector, win(), etc. by yourself!!! */

```

B Gnuplot

Gnuplot is a powerful tool for generating graphs. Here is a simple tutorial for using it for this assignment.

You should have your code stop training every thousand rounds and play 100 rounds against expecti-minimax. It should then spit out a file with the following format

```

1 5
2 7
3 10

```

In the above example, the output represents 5% win rate against expecti-minimax after 1000 rounds of training, and a 7% win rate after 2000 rounds.

Assuming that the above output is saved in a file called `bg_out.gnuplot.data`, we can create a control file for gnuplot generation. We will call this file `bg_out.gnuplot`.

```
set xlabel "% Win against expecti-minimax"
set ylabel "1000 self-training rounds"
set terminal postscript
set output "bg_out.ps"
set yrange [0:100]
plot "bg_out.gnuplot.data" notitle
```

You can then generate the postscript image (`bg_out.ps`) by calling the command:

```
gnuplot bg_out.gnuplot
```

Gnuplot is available on owl.net.