6.034 Quiz 3
14 November 2012

Name	
email	

Circle your TA **(for 1 extra credit point)**, so that we can more easily enter your score in our records and return your quiz to you promptly.

Dylan Holmes	**Sarah Lehmann**	**William Adams**
Igor Malioutov	**Robert McIntyre**	**Ami Patel**
Sila Sayan	**Mark Seifter**	**Stephen Serene**

Problem number	Maximum	Score	Grader
1	50		
2	50		
Total	100		

Problem number	Maximum	Score	Grader
3	6		

We recommend you reserve a few minutes for problem three.

There are 8 pages in this quiz, including this one, but not including blank pages and tear-off sheets. Tear-off sheets with duplicate drawings and data are located after the final page of the quiz. As always, open book, open notes, open just about everything, including a calculator, but no computers.

Problem 1: Neural Nets (50 points)

Part A (10 points)

In order to build a Terminator robot bodyguard, John Conner is designing a neural network that can recognize vertical edges. The inputs to the network consist of three pixels—a left pixel, a center pixel, and a right pixel—each of which can have a value between 0 (black) and 255 (white).

John decides that the **if the values of the left and right inputs differ by more than 100, then the network detects an edge, so the output of the network should be +1**. **Otherwise, the output of the network should be 0.**

Assign weights and thresholds to the network below to implement this behavior.

In the diagram below, the **boxes denote inputs**; all the **threshold inputs have a fixed value of -1**; and all neurons use the **unit step function** $step(x) = \begin{cases} 0, & x < 0 \\ 1, & x \geq 0 \end{cases}$, where *x* is the weighted sum of the neuron's inputs, and *step(x)* is the neuron's output.

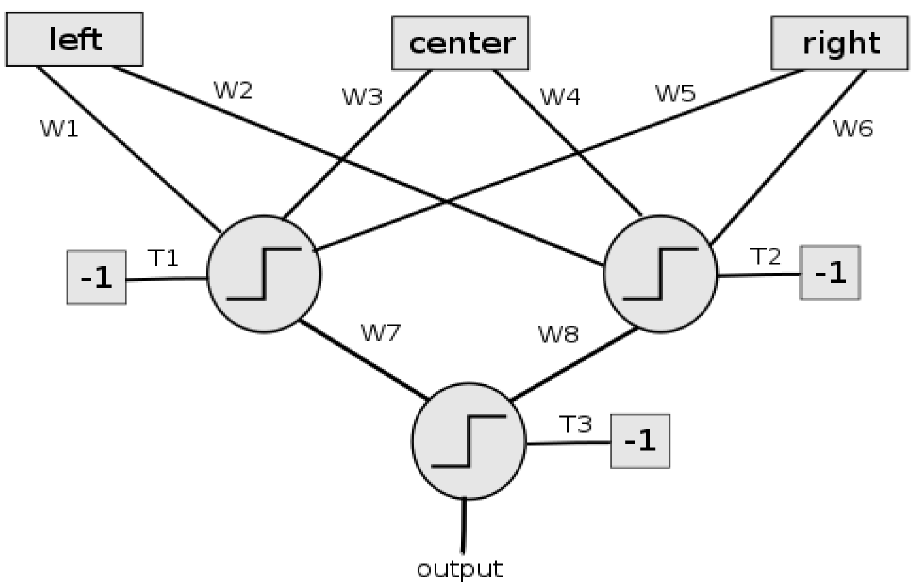

T1:	W1:	W5:
T2:	W2:	W6:
T3:	W3:	W7:
	W4:	W8:

Part B (40 points)

Now, John Conner has built a neural net which distinguishes human-shaped Terminator robots from real humans based on three input features: whether they **M**isuse colloquial expressions, whether they have **S**uper strength, and whether they are **G**ood with guns. An output value of **< ½ means that the network classifies the input as a human;** an output value of **> ½ means that the network classifies the input as a robot**.

Note: All neurons in this network use the **sigmoid function σ**(x) **with its threshold at zero** to compute their outputs.

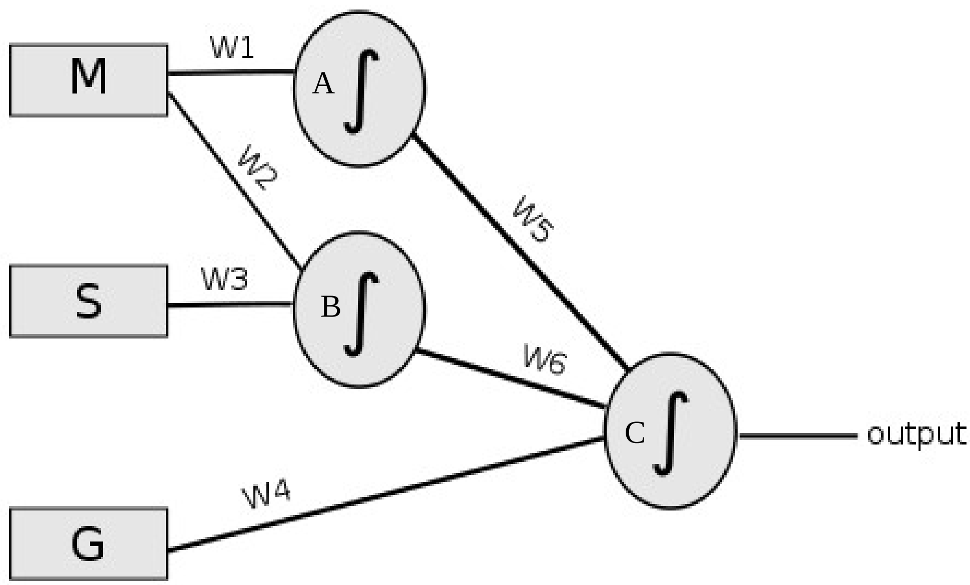

B1 (15 points): John decides to test this network on the character River from the show *Firefly.* River is a **human** (desired output = 0) with the following features:

Input Feature	**Value**
Misuses colloquial expressions (M)	1
Super strength (S)	0
Good with guns (G)	1

Write an expression for the network's output Z for River in terms of the input features, the unknown weights (W1, W2, etc.), and the sigmoid function σ(x).

B2 (15 points): Next, John Conner uses back propagation to train the neural net, using River as an example. **Write an expression for W1'**, the updated value of W1, in terms of the input features, the initial weights (W1, W2, etc.), the network output Z, and the training rate R. (Note: **Your expression might not include all of these variables.**)

B3 (5 points): How would the updated value for W1 compare to the value above if you started with a different initial value for W3 but kept all the other initial weights the same? **(Circle the best answer.)**

SAME DIFFERENT DEPENDS

Explain.

B4 (5 points): Now, John Connor performs a series of experiments using all sorts of networks, rates, and initial weights. In all his experiments, he trains his neural net with the same set of Terminator robots and humans (not including River). In some of the experiments, all the training examples are recognized correctly after training, but in those experiments, River is never classified correctly after training. What does this suggest about River?

Problem 2: Support Vector Machines (50 points)

Recently, you've been promised a lot of cake; unfortunately, many of those promises turn out to be false. Sometimes, you can get your hands on the supposed cakes' recipe beforehand. You decide to use this information to help **distinguish the cakes that are true (+) from the cakes that are lies (-)**.

You draw on prior experience to create a small training dataset for yourself, based on the cakes' sugar (x_s) and flour (x_f) content (shown below). Then, you apply SVMs to this problem.

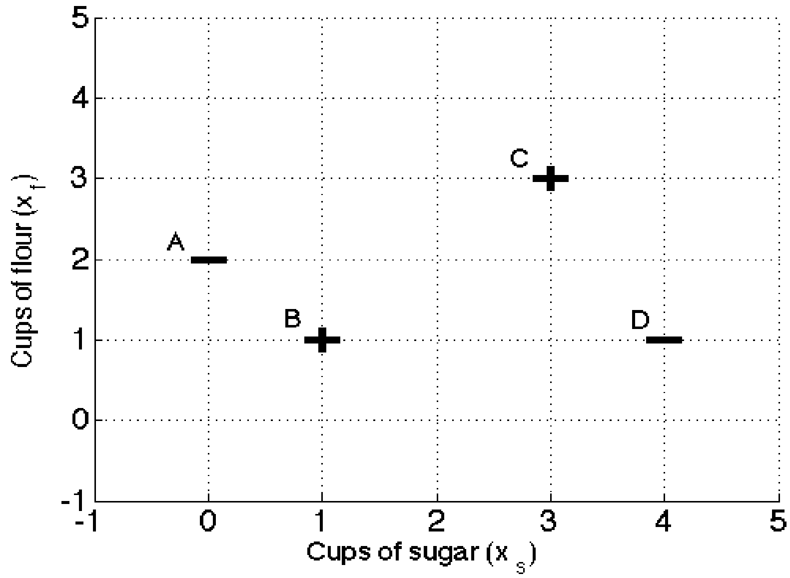

A (10 points)

Unfortunately your data is not linearly separable in terms of the features x_s and x_f. Atlas suggests using the following transformation on your feature vectors:

$$\varphi(x_s, x_f) = (|x_s - x_f|, x_f)$$

What is the corresponding kernel function K(**u**,**v**) in terms of u_s, u_f, v_s, and v_f where **u** = (u_s, u_f) and **v** = (v_s, v_f)?

B (15 points)

Draw the transformed training points on the following graph. **Label** the transformed training points with the same classification labels (+, -) and letters (A,B,C,D) as the original training points. **Draw** the SVM decision boundary with a heavy solid line and **draw** the gutters with heavy dashed lines. **Circle** the support vectors.

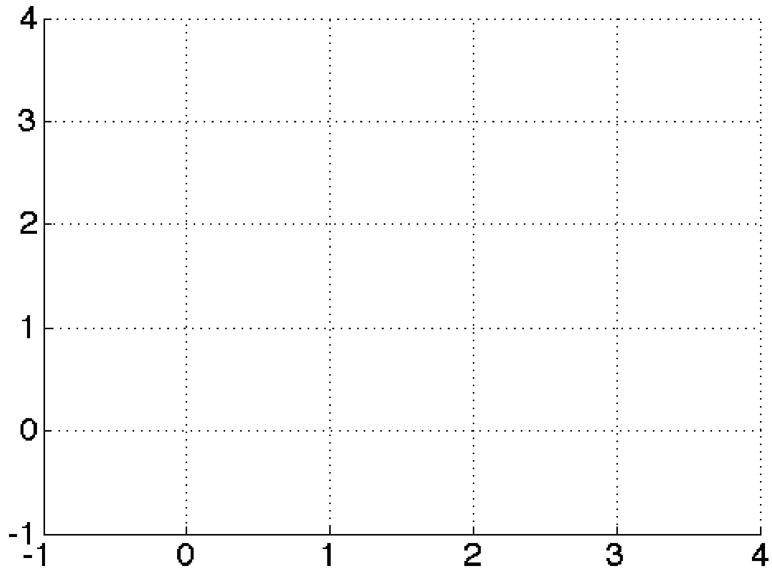

C (15 points)

Solve for the values of **w**, b, and the alphas for the training set.

w =

α_A =

α_C =

b =

α_B =

α_D =

Show work for partial credit:

D (5 points)

How would a recipe with 1⅔ cup granulated sugar and 2 cups all-purpose flour ($x_s = 1.\overline{6}$, $x_f = 2$) be classified by this SVM? **(Circle the best answer)**:

True (+) **Lie (-)** **Can't be determined**

Show work for partial credit:

E (5 points)

Someone promises you a cake with no flour ($x_f = 0$). For what values of sugar (x_s), will you classify such a cake as true (+)? (Include negative values of x_s, if any).

Answer: ____

Show work for partial credit:

Problem 3, Spiritual and Right Now

Circle the **best** answer for each of the following questions. There is **no penalty for wrong answers**, so it pays to guess in the absence of knowledge.

1 Boyden indicated it is now possible to:

1. Harness the power of cloud computing to accurately simulate mosquito brains.
2. Use ultrasonic resonance to shut off potassium ion transport across neuron cell walls .
3. Use Magnetic Resonance Imaging to measure changes in synaptic weights as learning occurs.
4. Embed photoreceptors in neuron cell walls.
5. Embed light-emitting diodes in neuron cell walls.

2 Kanwisher indicated we have special purpose brainware dedicated to processing images of:

1. Predators.
2. Body parts.
3. Insects.
4. Tools.
5. Paths.

3 Ullman indicated that:

1. Computers can easily recognize actions, such as drinking, even when performed by animals.
2. The computation required to recognize an object is exponential in the number object parts.
3. Knowledge of an object's parts can help separate instances of the object from background.
4. His recognition programs incorrectly identify scrambled collections of face parts as faces.
5. Separation of an object from background is the first and easiest part of object recognition.

4 Genetic algorithms can be viewed, in part, as a kind of:

1. Constraint propagation.
2. Hill climbing.
3. Resource allocation.
4. Support-vector machine.
5. Nearest-neighbor learning.

5 Learning phonological rules using the Sussman-Yip method is possible because

1. The number of examples is on the order of what is heard by a child in the first 3 years of life.
2. A language's phonemes densely pack distinctive feature space.
3. There are hundreds of distinctive features.
4. Negative examples prevent overgeneralization.
5. English does not exhibit the McGurk effect.

6 Near miss learning:

1. Is possible only if the teacher supplies positive examples only.
2. Is possible only if the teacher supplies large numbers of examples.
3. Is possible only if the teacher selects examples at random.
4. Makes it possible to learn something definite with every sample, positive or negative.
5. Combines the best features of neural nets and genetic algorithms.

Tear off page. You need not hand this page in.

For problem 1:

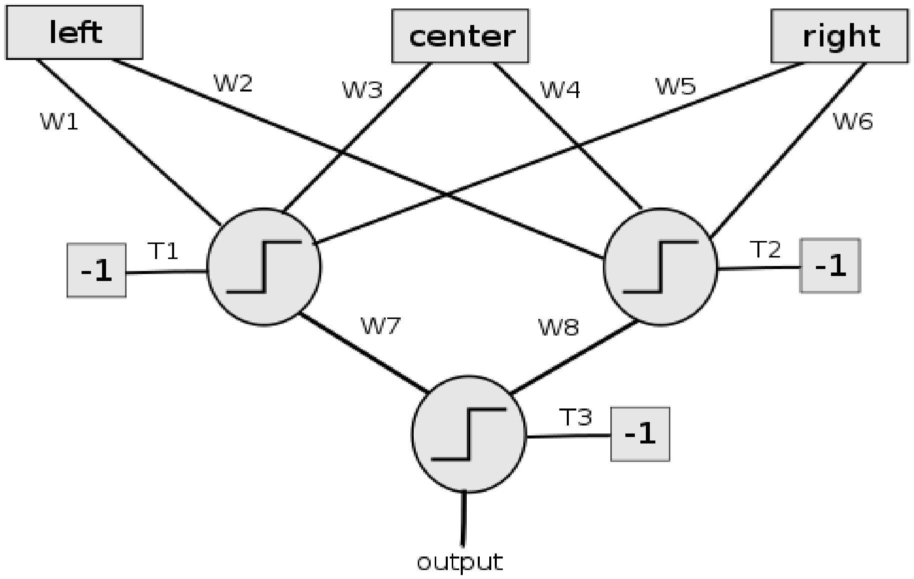

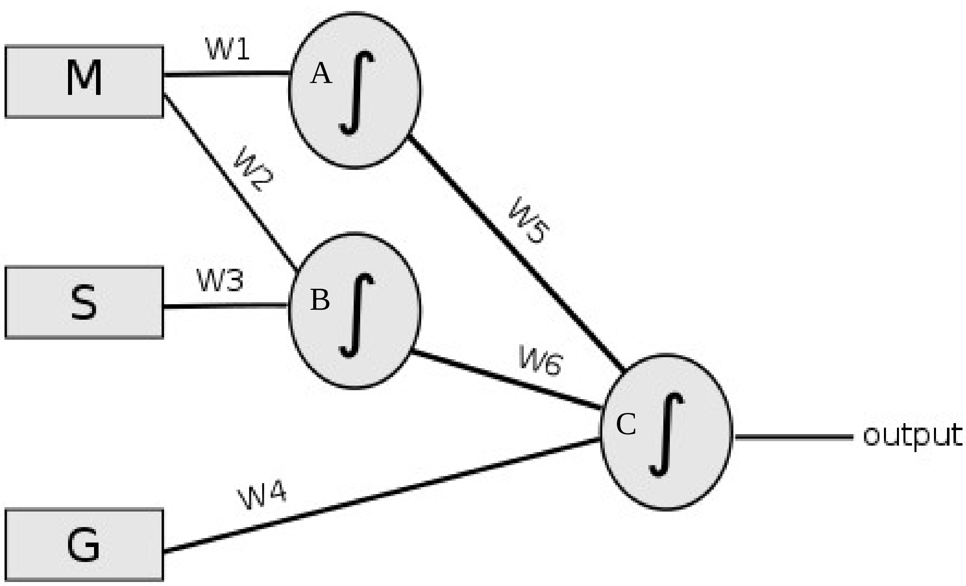

Tear off page. You need not hand this page in.

For problem 2: