CS/ENGRI 172, Fall 2003: Computation, Information, and Intelligence
9/26/03: Proof of PLA Convergence

More Useful Properties of the Inner Product

It should be fairly easy to convince yourself that for vectors $\vec{v}$, $\vec{w}$, $\vec{y}$, and $\vec{z}$ of the same dimensionality,

$$(\vec{v}+\vec{w}) \cdot (\vec{y}+\vec{z}) \quad = \quad \vec{v}\cdot\vec{y}+\vec{w}\cdot\vec{y}+\vec{v}\cdot\vec{z}+\vec{w}\cdot\vec{z}$$

For example, in the two-dimensional case,

$$\begin{aligned}
(\vec{v}+\vec{w}) \cdot (\vec{y}+\vec{z}) &= (v_1+w_1, v_2+w_2)\cdot(y_1+z_1, y_2+z_2) \\
&= (v_1+w_1)(y_1+z_1)+(v_2+w_2)(y_2+z_2) \\
&= (v_1y_1+w_1y_1+v_1z_1+w_1z_1)+(v_2y_2+w_2y_2+v_2z_2+w_2z_2) \\
&= (v_1y_1+v_2y_1)+(w_1y_1+w_2y_2)+(v_1z_1+v_2z_2)+(w_1z_1+w_2z_2) \\
&= \vec{v}\cdot\vec{y}+\vec{w}\cdot\vec{y}+\vec{v}\cdot\vec{z}+\vec{w}\cdot\vec{z}
\end{aligned}$$

This fact also implies, as a simpler case, that for any vectors $\vec{v}$, $\vec{y}$, and $\vec{z}$ of the same dimensionality,

$$\vec{v}\cdot(\vec{y}+\vec{z}) \quad = \quad \vec{v}\cdot\vec{y}+\vec{v}\cdot\vec{z}$$

Outline of our PLA Convergence Theorem Proof

We present a somewhat oblique, but therefore interesting, proof that the PLA converges. This is a general outline of the argument:

Given all of the constraints we have set for the oracle and the learner,

- First, define a *score* function which indicates how close the learner's weight vector $\vec{w}$ is to a "target"[1] vector $\vec{w}^*$. Our score measure will start at 0 and have a fractional form with a numerator and a denominator.
- Now show that every time the perceptron learning algorithm *updates*, so that its weight vector $\vec{w}_{old}$ is changed to $\vec{w}_{new}$, the score measure must increase by a non-negligible amount:
  - The score numerator must increase by *at least* g, the gap quantity.
  - The square of the score denominator must increase by *at most 1.*
- Therefore, after t updates, the score must have increased by at least $\sqrt{t}g$ from the initial score of 0.
- But, since our particular score measure has an upper-bound of 1, t can be *at most* $\frac{1}{g^2}$, which, since $g > 0$, shows that only a finite number of updates get made.

[1] Note the quotation marks! We are performing this proof from a vantage point outside of the oracle-learner system.