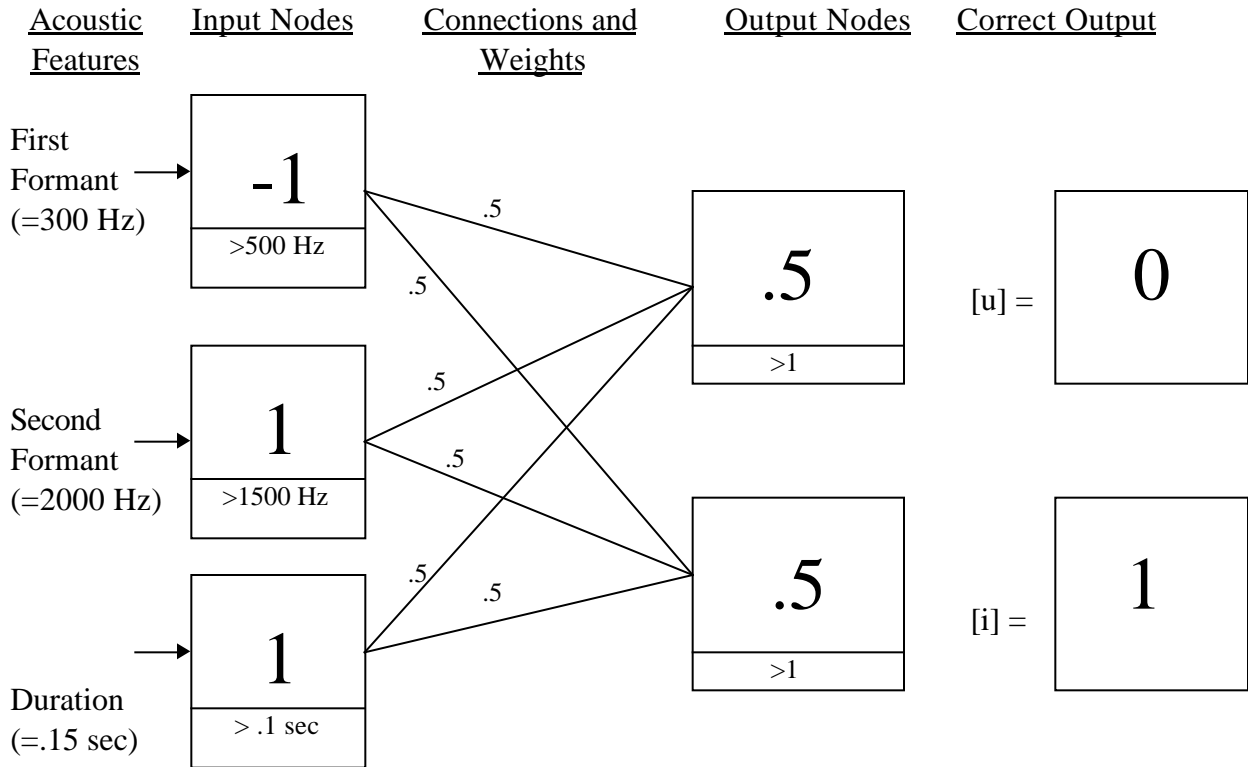


Training a Neural Net



What a neural net "knows" about the world is encoded in the connection weights between the nodes in its architecture. Neural nets are always born with a clean slate (as it were) to show how much knowledge-based behavior may arise from simply forming connections through experience and training. If a neural net could be trained to exhibit rule-based linguistic behavior, for example, this would argue against Chomsky's position that some knowledge of language must be innate.

Neural nets are "born knowing nothing" by setting all their original connection weights to some common value-- $.5$, in this example. The net will start "learning" things about the world by adjusting these weights in response to training--differently-valued weights will reflect what the network "knows" about the world.

Neural nets are trained by presenting them with stimuli and the correct categorizations of those stimuli. In this example, the network "hears" an [i] and is then told that it should have categorized that sound as an [i] (correct output values are 0 for [u] and 1 for [i]). Since the network's original connection weights were all the same, though, it yields activations that are identical for both [u] and [i]-- $.5$ each. The model will adjust its weights according to how far off its activations were from the correct values. A very simple adjustment might be half of the difference between the activation and correct values. Each connection to the [u] node, then, would decrease by $.25$ ($= (0-.5)/2$), while each connection to the [i] node would increase by $.25$ ($= (1-.5)/2$).

Training a Neural Net

Note: This training process would have to be repeated many times (usually in the hundreds or thousands) for the network to perceive natural speech reasonably well.