Subject: Re: slow processing of my k-nearest neighour code
Posted by JD Smith on Mon, 14 Aug 2006 21:39:18 GMT
View Forum Message <> Reply to Message

On Mon, 14 Aug 2006 10:11:30 -0400, Ben Tupper wrote:

> humphreymurray@gmail.com wrote:
>
>>> ; Calculate the squared distance for each attribute.
>>> squared = make_array(num_training_elements, num_attributes)
>>> for attrib = 0, num_attributes-1 do begin
>>> squared[*,attrib] = (testing_data[i, attrib] -
>>> training_data[*,attrib])^2
>>> endfor
>>>
>
> Hi,
>
> You might try replacing the above for inner-loop with the following
>
> squared = (testing_data - training_data)^2
>
> Since IDL is array saavy it will perform the operation element by
> element for you quite quickly (as well as make the "squared" array for you).
>
> You might be able to eliminate the outer-loop, too, but I am less sure
> of that. Take a peek at the for-loop bible at
>
> http://www.dfanning.com/tips/forloops.html

More of a diatribe than a bible ;).

I hope you meant:

http://www.dfanning.com/code_tips/slowloops.html

where the nearest-neighbor problem is actually discussed.

JD