On Thursday, March 7, 2013 11:12:35 AM UTC-7, David Fanning wrote:
> Paul van Delst writes:
>
>
>
>> If it makes you feel any better, I lay awake at night worrying about
>
>> plausibly incorrect results. In any language.
>
>
>
> Well, it *does* make me feel better. :-)
>
>
>
>> But, I don't really use IDL to produce quantitative results (computing
>
>> means and stddevs don't count)
>
>
>
> I was on a forum web page not too long ago and the folks there were
>
> going off on how horrible IDL was to do anything sensible and how it
>
> confounded expectations, was impossible for new people to learn, etc.
>
> Then some guy jumps in to say, 'Wait, there is a web site out there
>
> where the guy tries to fix all the things that are wrong with IDL.'
>
>
>
> It stopped me. Is that what I've been doing all these years? Fixing IDL?
>
>
>
> I suppose it is. And then I got totally depressed. Why do I do this
>
> work? And why do I do it for free, for God's sake!? I've no money in the
>
> bank, a job that pays very little, and I choose to spend my time fixing
>
> IDL. It's a joke.

>
>
>
> That's when I decided to go trekking in Patagonia. :-)
>
>
>
> Before I leave, though, I wrote my own histogram routine, cgHistogram,
>
> which I am just about to add to the Coyote Library. It fixes the bug
>
> with byte arrays (not fixed officially until IDL 8.2), checks to be sure
>
> the binsize data type matches the data type of the data you are binning,
>
> so you always get the correct results, allows for missing values in the
>
> histogram, will smooth the data if needed, and will return the relative
>
> frequency, instead of the histogram count, if you want that. It is a
>
> HELL of a lot better than the Histogram command, as always. Maybe one or
>
> two people will use it. ;-)
>
>
>
> Cheers,
>
>
>
> David
>
>
>
>
>
> --
>
> David Fanning, Ph.D.
>
> Fanning Software Consulting, Inc.
>
> Coyote's Guide to IDL Programming: http://www.idlcoyote.com/
>
> Sepore ma de ni thue. ("Perhaps thou speakest truth.")

It always amazed me IDL never did these things with the histogram function. It seems so obvious. I just imagined that every user has their own histogram wrapper function that does return the % rather than the count, and returns the actual x-axis of the histogram, etc.

cheers,
bob