Subject: string manipulation
Posted by Craig Markwardt on Mon, 26 Feb 2001 22:34:47 GMT
View Forum Message <> Reply to Message

I have been thinking about the following problem, but have been pretty frustrated.

I have input keyword/value strings of the following form. People may recognize this from a FITS file.

```
TTYPE2  = 'X1LSpecPcu0'        / X1LSpecPcu0 : Histogram
TUNIT2  = 'count   '
1CTYP2  = 'CHANNEL '
1CPIX2  = '0~4,5:53,(54~135;2),(136~237;3),(238~249;4),250~255'
12CD2A  = 1.25
^name^^   ^value^^^^^^^
```

All of the keyword names have a trailing "2" which indicates that they are describing column number 2 in a FITS table. Note that in the name "12CD2A", only the *final* 2 refers to the column number.

The interesting question happens when I want to change the column number, say from "2" to "50". Is there a straightforward way to do this in "vectoral" sort of way?

I am able to find the string positions of the 2's, so that's not really a problem. I do this by making a byte array of the strings, and blanking out any alphabetic characters and any leading numeric characters. Here I appreciate STRPOS is (partially) vectorized.

However, when it comes to resubstituting the "50" back in, that's when I get stymied. This is primarily because STRMID and STRPUT are not vectorized at all. Well STRMID *is* vectorized, but not with a sane behavior. For example, what I'd like to do is:

```
NEWKEY = STRMID(KEY,0,P1) + '50' + STRMID(KEY,P2,100)
```

Where KEY, P1, and P2 are vectors. Obviously this doesn't work. Any ideas?

Craig

--

Craig B. Markwardt, Ph.D. EMAIL: craigmnet@cow.physics.wisc.edu
Astrophysics, IDL, Finance, Derivatives | Remove "net" for better response

Subject: Re: string manipulation
Posted by landsman on Tue, 27 Feb 2001 21:29:00 GMT
View Forum Message <> Reply to Message

In article <on66hxm5tk.fsf@cow.physics.wisc.edu>, craigmnet@cow.physics.wisc.edu writes...
> This is primarily because STRMID and STRPUT are not
> vectorized at all. Well STRMID *is* vectorized, but not with a sane
> behavior. For example, what I'd like to do is:
>
> NEWKEY = STRMID(KEY,0,P1) + '50' + STRMID(KEY,P2,100)
> NEWKEY = STRMID(KEY,0,P1) + '50' + STRMID(KEY,P2,100)
>
> Where KEY, P1, and P2 are vectors. Obviously this doesn't work. Any
> ideas?

The problem when STRMID was vectorized for V5.3 was that it was made *too*
powerful -- it handles simultaneously both extraction from multiple strings
and multiple extractions from a single string. In practice, I think the
first situation -- extraction from multiple strings -- is far more common, but
has an ugly syntax in the current STRMID implementation. Here is how one
does the example above.

```
N = N_elements(KEY)
NEWKEY = STRMID(KEY,INTARR(1,N),REFORM(P1,1,N)) + '50' +  $
      STRMID(KEY,REFORM(P2,1,N), REPLICATE(100,1,N) )
```

I have thought about writing a simple wrapper around STRMID (say STRMIDV) that
would have a simpler syntax for the case of single extractions from
multiple strings.

--Wayne Landsman landsman@mpb.gsfc.nasa.gov

Subject: Re: string manipulation
Posted by John-David T. Smith on Tue, 27 Feb 2001 22:40:53 GMT
View Forum Message <> Reply to Message

Wayne Landsman wrote:
>
> In article <on66hxm5tk.fsf@cow.physics.wisc.edu>, craigmnet@cow.physics.wisc.edu writes...
>> This is primarily because STRMID and STRPUT are not
>> vectorized at all. Well STRMID *is* vectorized, but not with a sane
>> behavior. For example, what I'd like to do is:
>>
>> NEWKEY = STRMID(KEY,0,P1) + '50' + STRMID(KEY,P2,100)
>> NEWKEY = STRMID(KEY,0,P1) + '50' + STRMID(KEY,P2,100)
>>
>> Where KEY, P1, and P2 are vectors. Obviously this doesn't work. Any

>> ideas?
>
> The problem when STRMID was vectorized for V5.3 was that it was made *too*
> powerful -- it handles simultaneously both extraction from multiple strings
> and multiple extractions from a single string. In practice, I think the
> first situation -- extraction from multiple strings -- is far more common, but
> has an ugly syntax in the current STRMID implementation. Here is how one
> does the example above.
>
> N = N_elements(KEY)
> NEWKEY = STRMID(KEY,INTARR(1,N),REFORM(P1,1,N)) + '50' + $
> STRMID(KEY,REFORM(P2,1,N), REPLICATE(100,1,N))
>
> I have thought about writing a simple wrapper around STRMID (say STRMIDV) that
> would have a simpler syntax for the case of single extractions from
> multiple strings.

Aha! Where there's a will... Pavel, rejoin your faith.

A bit of a refinement, for the lazy among us:

```
new=strmid(key,0,transpose(p1))+'50'+strmid(key,tranpose(p2) )
```

The key is putting the threading vector on its head, as a column vector. Another simplification arrives from strmid's willingness to loop back over vectors which are too short (like the scalar 0), and to extract all the way to the end of a string, if no length is specified.

Row vectors are interpreted as multiple places in each given string to operate. Power, with a price.

Note that an even easier notation appears if you have p1 and p2 as columns in an array, e.g.:

```
IDL> p=[ [1,5], [2,6], [3,6], [4,8] ]
```

then you can simply use the relatively clean:

```
new=strmid(key,0,p[0,*])+'50'+strmid(key,p[1,*])
```

Here is a perfect case of where IDL's notion of keeping leading dimensions of size 1 is critical. Note that p[0,*] are lengths, and p[1,*] are subscripts.

This also works quite well for p1 and p2 not the same length as key, or each other.

JD

Subject: Re: string manipulation
Posted by Mark Hadfield on Tue, 27 Feb 2001 23:02:00 GMT
View Forum Message <> Reply to Message

"JD Smith" <jdsmith@astro.cornell.edu> wrote in message
news:3A9C2CF5.F632B1AF@astro.cornell.edu...
> [...]
> new=strmid(key,0,transpose(p1))+'50'+strmid(key,tranpose(p2))
>
> The key is putting the threading vector on its head, as a column
> vector.

Aaaaagh! The major reason I hate Matlab is all the "Have I got a row vector or a column vector?" rubbish! I feel like picking it up & shaking it & saying, "No, it's not a matrix, it's not a column vector, it's not a row vector, it's just a list of numbers. It's not dimensioned (n,1) and it's not dimensioned (1,n), it's dimensioned (n). You can print it vertically, you can print it horizontally, I don't care, it's just a list of numbers!"

And now JD shows us how to confuse ourselves with row vectors and column vectors in IDL! As if HISTOGRAM wasn't enough. JD, I call down a curse on you, and your children, and your children's chil.....[falls from chair in paroxysm of rage]

Mark Hadfield
m.hadfield@niwa.cri.nz http://katipo.niwa.cri.nz/~hadfield
National Institute for Water and Atmospheric Research

Subject: Re: string manipulation
Posted by davidf on Tue, 27 Feb 2001 23:34:44 GMT
View Forum Message <> Reply to Message

Mark Hadfield (m.hadfield@niwa.cri.nz) writes:

> And now JD shows us how to confuse ourselves with row vectors and column
> vectors in IDL! As if HISTOGRAM wasn't enough. JD, I call down a curse on
> you, and your children, and your children's chil.....[falls from chair in
> paroxysm of rage]

Poor Mark. He's having a really bad day. Does
anyone know the URL of that place where you
can send a virtual bouquet of flowers to someone ...

Cheers,

David

--
David Fanning, Ph.D.
Fanning Software Consulting
Phone: 970-221-0438 E-Mail: davidf@dfanning.com
Coyote's Guide to IDL Programming: http://www.dfanning.com/
Toll-Free IDL Book Orders: 1-888-461-0155

Subject: Re: string manipulation
Posted by John-David T. Smith on Wed, 28 Feb 2001 00:35:35 GMT
View Forum Message <> Reply to Message

Mark Hadfield wrote:
>
> "JD Smith" <jdsmith@astro.cornell.edu> wrote in message
> news:3A9C2CF5.F632B1AF@astro.cornell.edu...
>> [...]
>> new=strmid(key,0,transpose(p1))+'50'+strmid(key,tranpose(p2))
>>
>> The key is putting the threading vector on its head, as a column
>> vector.
>
> Aaaaagh! The major reason I hate Matlab is all the "Have I got a row vector
> or a column vector?" rubbish! I feel like picking it up & shaking it &
> saying, "No, it's not a matrix, it's not a column vector, it's not a row
> vector, it's just a list of numbers. It's not dimensioned (n,1) and it's not
> dimensioned (1,n), it's dimensioned (n). You can print it vertically, you
> can print it horizontally, I don't care, it's just a list of numbers!"
>
> And now JD shows us how to confuse ourselves with row vectors and column
> vectors in IDL! As if HISTOGRAM wasn't enough. JD, I call down a curse on
> you, and your children, and your children's chil.....[falls from chair in
> paroxysm of rage]

Good on ya. I aim to please.

JD

P.S. The many ways to make column vectors:

IDL> cv=[[1],[2],[3]]

or

IDL> cv=rotate([1,2,3],1)

or

IDL> cv=transpose([1,2,3])

or

IDL> cv=array[0,*]

or

IDL> cv=reform([1,2,3],1,3)

or

IDL> cv=[1,2,3]##1

or

IDL> cv=1#[1,2,3]

or

...

Subject: Re: string manipulation
Posted by Mark Hadfield on Wed, 28 Feb 2001 20:30:52 GMT
View Forum Message <> Reply to Message

"David Fanning" <davidf@dfanning.com> wrote in message
news:MPG.1505e77da66a754e989d6f@news.frii.com...
>
> Poor Mark. He's having a really bad day. Does
> anyone know the URL of that place where you
> can send a virtual bouquet of flowers to someone ...

Nobody's taken the hint yet!

Does anyone know the URL of that place where you can send a virtual pint of
beer to someone ...

Mark Hadfield
m.hadfield@niwa.cri.nz http://katipo.niwa.cri.nz/~hadfield
National Institute for Water and Atmospheric Research

Subject: Re: string manipulation

Posted by Kenneth P. Bowman on Thu, 01 Mar 2001 00:16:23 GMT
View Forum Message <> Reply to Message

In article <983392306.931105@clam-ext>, Mark Hadfield
<m.hadfield@niwa.cri.nz> wrote:

> Does anyone know the URL of that place where you can send a virtual pint of
> beer to someone ...

No, but I expect someone could whip up an object to *render* a virtual
pint of beer. But wait, would that be hardware or software rendering?

Ken

Subject: Re: string manipulation
Posted by Pavel A. Romashkin on Thu, 01 Mar 2001 19:54:36 GMT
View Forum Message <> Reply to Message

"Kenneth P. Bowman" wrote:
>
> No, but I expect someone could whip up an object to *render* a virtual
> pint of beer. But wait, would that be hardware or software rendering?

Just make sure it works on my Mac... and does not render a Budweiser. I
don't care whether it uses OpenGL to open the bottle or a hardware tool.

Pavel
