Subject: Wiener filter
Posted by Richard Tyc on Mon, 17 Dec 2001 20:41:40 GMT
View Forum Message <> Reply to Message

Has anyone developed a Wiener filter algorithm for image processing in IDL
(and be willing to share it ???)
My image processing handbook by John Russ does not have it ??
A paper that describes it says it produces a "minimum least-squares error
between the "true" uncorrupted image and the noisy, measured version" It
makes use of the power spectral density of the image.

Any help appreciated.....

Thanks
Rich

Subject: Re: Wiener filter
Posted by jeyadev on Tue, 18 Dec 2001 20:30:54 GMT
View Forum Message <> Reply to Message

In article <9vllbg$po7$1@canopus.cc.umanitoba.ca>,
Richard Tyc <richt@sbrc.umanitoba.ca> wrote:
> Has anyone developed a Wiener filter algorithm for image processing in IDL
> (and be willing to share it ???)
> My image processing handbook by John Russ does not have it ??
> A paper that describes it says it produces a "minimum least-squares error
> between the "true" uncorrupted image and the noisy, measured version" It
> makes use of the power spectral density of the image.
>
> Any help appreciated.....

I am note sure what you mean by 'filter'. The Wiener spectrum is
defined by precisely what you quote: it is the square of the
Fourier transform of the (density) fluctuations.

Let D(x,y) be the initial data in real space, i.e. your image.
Let the average density be denoted by <D>.

Let d(p,q) be the Fourier transform of D(x,y) - <D>, i.e.

$$d(p,q) = FT\{ D(x,y) - \langle D\rangle \}$$

Then, the Wiener spectrum is given by

$$W(p,q) = |d(p,q)|^2 = | FT\{ D(x,y) - \langle D\rangle \} |^2$$

I do not see how this can be viewed as a filter. It is merely the

square of the amplitudes of the Fourier transform of the fluctuations
from the mean.

Am I missing something in your question?
--

Surendar Jeyadev jeyadev@wrc.xerox.com

Subject: Re: Wiener filter
Posted by James Kuyper Jr. on Wed, 19 Dec 2001 15:03:21 GMT
View Forum Message <> Reply to Message

Surendar Jeyadev wrote:

> In article <9vllbg$po7$1@canopus.cc.umanitoba.ca>,
> Richard Tyc <richt@sbrc.umanitoba.ca> wrote:
>
>> Has anyone developed a Wiener filter algorithm for image processing in IDL
>> (and be willing to share it ???)
>> My image processing handbook by John Russ does not have it ??
>> A paper that describes it says it produces a "minimum least-squares error
>> between the "true" uncorrupted image and the noisy, measured version" It
>> makes use of the power spectral density of the image.
>>
>> Any help appreciated.....
>
>
> I am note sure what you mean by 'filter'. The Wiener spectrum is
> defined by precisely what you quote: it is the square of the
> Fourier transform of the (density) fluctuations.
...

Optimal Wiener filtering of a one-dimensional data set is described in
section 12.6 of "Numerical Recipes in C", by Preuss et.al. It cites
three books on signal processing as references. The basic result is that
if you have a corrupted signal with the fourier spectrum S(f),
containing noise with a fourier specturm N(f), it can be shown
rigorously that the optimal (in the sense of a least-squares fit)
frequency filter for removing the noise is:

```
             |S(f)|^2
 phi(f) = --------------------
        |S(f)|^2 + |N(f)|^2
```

The procedure is straightforward. Estimate the fourier spectrum of the
noise. Calculate the fourier spectrum of the corrupted signal. Calculate
the corresponding filter function. Multiply the fourier spectrum of the

corrupted signal by the filter function. Do an inverse fourier transform
on the resulting function, to get an optimum estimate to the uncorrupted
signal.

A seemingly bright idea is to subtract that estimate of the uncorrupted
signal from the corrupted signal, to get a better estimate of the noise.
Then use that improved noise estimate to re-filter the data. This
doesn't work: the process converges to a signal of S(f)=0. Basically,
each step in that process treats a portion of the uncorrupted signal as
if it were part of the noise.

I've not seen this derived for an image processing context, where you
have a two-dimension fourier transform to contend with, but I would
expect the results to be basically the same.

Subject: Re: Wiener filter
Posted by Richard Tyc on Wed, 19 Dec 2001 20:10:29 GMT
View Forum Message <> Reply to Message

```
> Optimal Wiener filtering of a one-dimensional data set is described in
> section 12.6 of "Numerical Recipes in C", by Preuss et.al. It cites
> three books on signal processing as references. The basic result is that
> if you have a corrupted signal with the fourier spectrum S(f),
> containing noise with a fourier specturm N(f), it can be shown
> rigorously that the optimal (in the sense of a least-squares fit)
> frequency filter for removing the noise is:
>
>                |S(f)|^2
> phi(f) = --------------------
>          |S(f)|^2 + |N(f)|^2
>
```

What exactly is |S(f)|^2

If I have a 2D corrupted image, say I(x,y)

Is it   ABS(FFT(I)) ^ 2 or the magnitude of the complex FFT result
squared (Power Spectrum) squared ?

```
> The procedure is straightforward. Estimate the fourier spectrum of the
> noise. Calculate the fourier spectrum of the corrupted signal. Calculate
> the corresponding filter function. Multiply the fourier spectrum of the
> corrupted signal by the filter function. Do an inverse fourier transform
> on the resulting function, to get an optimum estimate to the uncorrupted
> signal.
```

It seems some knowledge of the noise is required. What if it was modeled as 'white noise' where it would be constant at all spatial frequencies.

A paper I am using that discusses this in the context of my problem points out, "....Assuming that noise power spectrum is white, the mean spectral density at high spatial frequencies was calculated and subtracted from P(f) (the power spectral density of the corrupted image) to estimate S(f) (power spectral density of uncorrupted image). Can you shed any light on this in terms of IDL code ??

Thanks for the help
Rich

Subject: Re: Wiener filter
Posted by jeyadev on Wed, 19 Dec 2001 21:34:01 GMT
View Forum Message <> Reply to Message

In article <3C20AC39.2080907@gsfc.nasa.gov>,
James Kuyper Jr. <James.R.Kuyper.1@gsfc.nasa.gov> wrote:
> Surendar Jeyadev wrote:
>
>> In article <9vllbg$po7$1@canopus.cc.umanitoba.ca>,
>> Richard Tyc <richt@sbrc.umanitoba.ca> wrote:
>
> Optimal Wiener filtering of a one-dimensional data set is described in
> section 12.6 of "Numerical Recipes in C", by Preuss et.al. It cites
> three books on signal processing as references. The basic result is that
> if you have a corrupted signal with the fourier spectrum S(f),
> containing noise with a fourier specturm N(f), it can be shown
> rigorously that the optimal (in the sense of a least-squares fit)
> frequency filter for removing the noise is:
>
```
>              |S(f)|^2
>  phi(f) = -------------------
>        |S(f)|^2 + |N(f)|^2
```

Eeeks! I did not know the name. I have actually used it is some data analysis in 1-d (not 2-d images). Thanks!
--

Surendar Jeyadev   jeyadev@wrc.xerox.com

Subject: Re: Wiener filter
Posted by James Kuyper Jr. on Wed, 19 Dec 2001 21:48:57 GMT
View Forum Message <> Reply to Message

Richard Tyc wrote:

>> Optimal Wiener filtering of a one-dimensional data set is described in
>> section 12.6 of "Numerical Recipes in C", by Preuss et.al. It cites
>> three books on signal processing as references. The basic result is that
>> if you have a corrupted signal with the fourier spectrum S(f),
>> containing noise with a fourier specturm N(f), it can be shown
>> rigorously that the optimal (in the sense of a least-squares fit)
>> frequency filter for removing the noise is:
>>

```
>>               |S(f)|^2
>>  phi(f) = --------------------
>>          |S(f)|^2 + |N(f)|^2
>>
```

>
>
> What exactly is |S(f)|^2

It's proportional to the power spectral density of the corrupted signal.
S(f) is the fourier transform of the the corrupted signal.

> If I have a 2D corrupted image, say I(x,y)
>
> Is it ABS(FFT(I)) ^ 2 or the magnitude of the complex FFT result
> squared (Power Spectrum) squared ?

It's ABS(FFT(I))^2. I think that the following code should be a more
efficient way to calculate the same value:

```
s = FFT(I)
ps = DOUBLE(s,conj(s))
```

...

> It seems some knowledge of the noise is required. ...

Correct.

> ... What if it was modeled as
> 'white noise' where it would be constant at all spatial frequencies.

Optimal Wiener filtering is a way of using your knowledge of the
frequency power spectrum of the noise, to extract it from the data. If
you're using a "white noise" spectrum because you know that your noise
has that characteristic, that's reasonable. If you're using "white
noise" because you're not sure what the noise spectrum looks like, and
are afraid to commit yourself, Wiener filtering is inappropriate. Keep

in mind that you need to know not merely the frequency dependence of the noise, but also it's absolute magnitude.

On the other hand, you don't need to know the noise power spectrum very precisely. The result of the filtering is insensitive to small errors in the assumed noise spectrum, just like sqrt(x) is insensitve to small errors in 'x'.

> A paper I am using that discusses this in the context of my problem points
> out, "....Assuming that noise power spectrum is white, the mean spectral
> density at high spatial frequencies was calculated and subtracted from P(f)
> (the power spectral density of the corrupted image) to estimate S(f) (power
> spectral density of uncorrupted image). Can you shed any light on this in
> terms of IDL code ??

Note a difference in notation here: that quotation identifies S(f) as the power spectral density of the uncorrupted image. I was using that same notation to mean the fourier transform of the corrupted signal.

I'm afraid I can't convert that to code; the only tricky step is one that they've given no details about. That step is the one where they estimate the power spectrum of the noise. They said that they were assuming a "white noise" spectrum, but that still leaves them with the problem of estimating the amplitude of the noise. One plausible approach is to plot the power spectrum of your signal, and decide to model it as the sum of two simple curves with known shapes. Then use regress() to fit the data to a linear sum of those two curves.

The part they do explain is trivial. Using the notation from that quote, P(f) = S(f) - N(f), where P(f) is the power spectrum of the uncorrupted signal, S(f) is the power specttrum of the corrupted signal, and N(f) is the power spectrum of the noise. (Note change of notation from previous context).

Subject: Re: Wiener filter
Posted by James Kuyper Jr. on Wed, 19 Dec 2001 21:50:42 GMT
View Forum Message <> Reply to Message

James Kuyper Jr. wrote:
...

> ps = DOUBLE(s,conj(s))

Correction:
ps = DOUBLE(s*conj(s))

Subject: Re: Wiener filter
Posted by Richard Tyc on Fri, 21 Dec 2001 18:40:11 GMT
View Forum Message <> Reply to Message

> I'm afraid I can't convert that to code; the only tricky step is one
> that they've given no details about. That step is the one where they
> estimate the power spectrum of the noise. They said that they were
> assuming a "white noise" spectrum, but that still leaves them with the
> problem of estimating the amplitude of the noise. One plausible approach
> is to plot the power spectrum of your signal, and decide to model it as
> the sum of two simple curves with known shapes. Then use regress() to
> fit the data to a linear sum of those two curves.
>
> The part they do explain is trivial. Using the notation from that quote,
> P(f) = S(f) - N(f), where P(f) is the power spectrum of the uncorrupted
> signal, S(f) is the power specttrum of the corrupted signal, and N(f) is
> the power spectrum of the noise. (Note change of notation from previous
> context).

Thanks for your help. I now understand the process a little better but I too
am still unclear on the noise amplitude estimation. I don't quite follow
your idea of "sum of two separate curves" and then using regress().

I have stumbled into a fairly sophisticated subject her. Could you point me
to some references that may explain your idea in more detail ?
The paper does refer to : "Digital Image Processing" by Gonzales which I
have on order AND "Numerical Recipes: the art of scientific computing" by
Press, Flannery et al which I should be able to find around here.
Any others?

Rich

Subject: Re: Wiener filter
Posted by James Kuyper Jr. on Fri, 21 Dec 2001 20:26:27 GMT
View Forum Message <> Reply to Message

Richard Tyc wrote:
...

> Thanks for your help. I now understand the process a little better but I too
> am still unclear on the noise amplitude estimation. I don't quite follow
> your idea of "sum of two separate curves" and then using regress().
>
> I have stumbled into a fairly sophisticated subject her. Could you point me
> to some references that may explain your idea in more detail ?
> The paper does refer to : "Digital Image Processing" by Gonzales which I
> have on order AND "Numerical Recipes: the art of scientific computing" by

> Press, Flannery et al which I should be able to find around here.
> Any others?

Well, everything I've ever read about the subject is in "Numerical Recipes". If you can find that, I can't give you any better citations. The book itself contains three citations in that section which you could follow up on. That idea of "sum of two seperate curves" is explained graphically in Figure 12.6.1. Whether or not you can use regress() depends upon whether the combined curve is linear in the unknown parameters of the individual curves. If it's non-linear in the unknown parameters, you'll have to use more sophisticated fitting techniques.

Note: I've apparantly made a notational error while explaining this. S(f) is the fourier transform of signal you want to extract; I'd been implying that it was the fourier transform of the signal plus the noise. He actually uses C(f) for that purpose. I'm sorry for causing any confusion! This shows you how often I've actually used this technique! I know a lot of things about a lot of nifty numerical techniques that I've never been able to put to actual use. :-(
