Subject: Re: Averaging quaternions

Posted by Craig Markwardt on Thu, 18 Mar 2004 23:39:43 GMT

View Forum Message <> Reply to Message

GrahamWilsonCA@yahoo.ca (Graham) writes:

- > Does anyone know if it is possible to take an average of regularly
- > sampled quaternions to get a mean orientation (i.e. a mean rotation
- > matrix)? I seem to recall there being a trick involved but beyond
- > re-normalizing the resuling (averaged) quaternion, I cannot remember
- > what it is.

I am sure I will be scolded by somebody, but I believe that you can average the quaternion components, and then normalize as you say. This is assumes that you are noise dominated.

Also, there is one trick that I can think of, which is that quaternions are degenerate. For each unique rotation, there are two possible quaternions whose components have opposite signs. This is because a positive rotation about axis V is identical to a negative rotation about axis -V.

If your system is capable of both signs indiscriminately, then you must make the sign conventions uniform. For example, by always making one component positive.

Some advertising since you crossposted on the IDL newsgroup: I do have a fairly comprehensive quaternion IDL library on my web page.

Craig

P.S.	http://cow.phys	sics.wisc.edu/~craigm/idl/idl.html	(under Math)

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

Subject: Re: Averaging quaternions
Posted by One Usenet Poster on Fri, 19 Mar 2004 04:17:40 GMT
View Forum Message <> Reply to Message

Graham wrote:

- > Does anyone know if it is possible to take an average of regularly
- > sampled quaternions to get a mean orientation (i.e. a mean rotation

- > matrix)? I seem to recall there being a trick involved but beyond
- > re-normalizing the resuling (averaged) quaternion, I cannot remember
- > what it is.

>

- > Cheers,
- > Graham

Graham:

Maybe you could try the SLERP (Spherical Linear IntERPolation) method to average the quaternions. Here's a link that describes the SLERP method: http://www.theory.org/software/qfa/writeup/node12.html

Good luck, OUP

Subject: Re: Averaging quaternions
Posted by John Lansberry on Fri, 19 Mar 2004 14:53:45 GMT
View Forum Message <> Reply to Message

"Craig Markwardt" <craigmnet@REMOVEcow.physics.wisc.edu> wrote in message news:on65d167y8.fsf@cow.physics.wisc.edu...

> GrahamWilsonCA@yahoo.ca (Graham) writes:

>

- >> Does anyone know if it is possible to take an average of regularly
- >> sampled quaternions to get a mean orientation (i.e. a mean rotation
- >> matrix)? I seem to recall there being a trick involved but beyond
- >> re-normalizing the resuling (averaged) quaternion, I cannot remember
- >> what it is.

>

- > I am sure I will be scolded by somebody, but I believe that you can
- > average the quaternion components, and then normalize as you say.
- > This is assumes that you are noise dominated.

>

Averaging components is a bad idea no matter what, since the result is never a "quaternion." The OP doesn't imply anything about "noise."

- > Also, there is one trick that I can think of, which is that
- > quaternions are degenerate. For each unique rotation, there are two
- > possible quaternions whose components have opposite signs. This is
- > because a positive rotation about axis V is identical to a negative
- > rotation about axis -V.

>

- > If your system is capable of both signs indiscriminately, then you
- > must make the sign conventions uniform. For example, by always making
- > one component positive.

You are correct that q and -q represent the same rotation - that's not "degenerate", it's just not "unique." Typically, the "scalar" part of the quaternion, cos(theta/2), is chosen to be the component that's always positive.

John

Subject: Re: Averaging quaternions
Posted by John Lansberry on Fri, 19 Mar 2004 15:14:12 GMT
View Forum Message <> Reply to Message

"Graham" < GrahamWilsonCA@yahoo.ca> wrote in message news:eda30d78.0403181434.229b3b53@posting.google.com...

- > Does anyone know if it is possible to take an average of regularly
- > sampled quaternions to get a mean orientation (i.e. a mean rotation
- > matrix)? I seem to recall there being a trick involved but beyond
- > re-normalizing the resuling (averaged) quaternion, I cannot remember
- > what it is.

>

- > Cheers.
- > Graham

Hi Graham,

It might help a lot to explain a little more about what you mean by "regularly sampled" and "averaging."

Do you mean that you have time samples of quaternions over time? In other words, you have q(t1), q(t2), q(t3), etc., where t1, t2, t3, etc. are evenly spaced time points (not that even spacing is all that important).

Is each sample a valid quaternion?

By "averaging," what do you really want?

SLERP (as suggested by OUP) is generally used to interpolate between two (valid) quaternions. The result is always a valid quaternion. SLERP is in fact valid for any two unit vectors of equal dimension (not just quaternions). So, you can use SLERP to find the quaternion at the midpoint of your time interval and call that an "average".

My guess is that you need something more complex - try Googling on "averaging quaternions".

John

Subject: Re: Averaging quaternions
Posted by John Lansberry on Fri, 19 Mar 2004 15:22:31 GMT
View Forum Message <> Reply to Message

Sorry, I didn't finish before sending. I should have mentioned, however, that Craig's suggestion to just "average the components and normalize" is, in fact, a common approach (see, for example, CLAUS GRAMKOW, "On Averaging Rotations", International Journal of Computer Vision 42(1/2), 7-16, 2001).

So Craig's suggestion is certainly one method (just not one I happen to like very much).

No scolding intended.

John

>

>>

>>

>>

>>

"John Lansberry" <john.lansberry@jhuapl.edu> wrote in message news:c3f1hp\$fto\$1@aplcore.jhuapl.edu...

> "Craig Markwardt" <craigmnet@REMOVEcow.physics.wisc.edu> wrote in message

> news:on65d167y8.fsf@cow.physics.wisc.edu...

>> GrahamWilsonCA@yahoo.ca (Graham) writes:

>>> Does anyone know if it is possible to take an average of regularly

>>> sampled quaternions to get a mean orientation (i.e. a mean rotation

>>> matrix)? I seem to recall there being a trick involved but beyond

>>> re-normalizing the resuling (averaged) quaternion, I cannot remember

>>> what it is.

>> I am sure I will be scolded by somebody, but I believe that you can

>> average the quaternion components, and then normalize as you say.

>> This is assumes that you are noise dominated.

> Averaging components is a bad idea no matter what, since the result is never

> a "quaternion." The OP doesn't imply anything about "noise."

>> Also, there is one trick that I can think of, which is that

>> quaternions are degenerate. For each unique rotation, there are two

>> possible quaternions whose components have opposite signs. This is

>> because a positive rotation about axis V is identical to a negative

>> rotation about axis -V.

>> If your system is capable of both signs indiscriminately, then you

>> must make the sign conventions uniform. For example, by always making

>> one component positive.

> You are correct that q and -q represent the same rotation - that's not

- "degenerate", it's just not "unique." Typically, the "scalar" part of thequaternion, cos(theta/2), is chosen to be the component that's always
- >
- > John

> positive.

- >
- >

Subject: Re: Averaging quaternions
Posted by Arnold Neumaier on Fri, 19 Mar 2004 16:24:26 GMT
View Forum Message <> Reply to Message

John Lansberry wrote:

- > "Craig Markwardt" <craigmnet@REMOVEcow.physics.wisc.edu> wrote in message
- > news:on65d167y8.fsf@cow.physics.wisc.edu...

>

- >> GrahamWilsonCA@yahoo.ca (Graham) writes:
- >>> Does anyone know if it is possible to take an average of regularly
- >>> sampled quaternions to get a mean orientation (i.e. a mean rotation
- >>> matrix)? I seem to recall there being a trick involved but beyond
- >>> re-normalizing the resuling (averaged) quaternion, I cannot remember
- >>> what it is.

>

- >> quaternions are degenerate. For each unique rotation, there are two
- >> possible quaternions whose components have opposite signs. This is
- >> because a positive rotation about axis V is identical to a negative
- >> rotation about axis -V.

>>

- >> If your system is capable of both signs indiscriminately, then you
- >> must make the sign conventions uniform. For example, by always making
- >> one component positive.

This will not work if the component is zero. Averaging (0,q) and (0,-q) which are the same rotation gives 0, which is meaningless.

Thus the average and scale procedure makes only sense if all quaternions are oriented the same way. One way to achieve this for any set of unit quaternions that do not stray too much is the following:

- 1. apply to all quaternions a rotation that moves one of them to 1 (for example one that is closest to the trivial average),
- 2. orient all results to positive real part,
- 3. average the results,
- 4. rotate back the result,

5. normalize.

More costly but completely reliable.

Arnold Neumaier

Subject: Re: Averaging quaternions
Posted by GrahamWilsonCA on Mon, 29 Mar 2004 20:45:02 GMT
View Forum Message <> Reply to Message

I have now made a few attempts at averaging my quaternions. For what it is worth, the quaternion data I have is from a star-tracker on board a satellite and each quaternion represents the inferred spacecraft attitude from an algebraic computation using the positions of 3-6 stars on a CCD.

The paper by Curtis et al. in the 1993 IEEE Annual virtual reality international symposium was probably the easiest for me to actually understand so I started with it. However, their approach only works for averaging 2 quaternions and I have 15-1000 to average in time. Moreover, I need to weight each quaternion by its relative precision (i.e. the number of stars on the CCD and the amount of stray light) at any given time. Both of these points are eluted to in the future work section of this paper but I cannot find any further papers by these people.

From the suggested approach:

- 1. apply to all quaternions a rotation that moves one of them to 1(for example one that is closest to the trivial average),
- > 2. orient all results to positive real part,
- > 3. average the results
- > 4. rotate back the result,
- > 5. normalize.

It isn't clear why steps 1 and 4 are required as they can be combined with 2 using a dot product? I'm not entirely sure of that but the following seems to work ok:

```
Quaternion_t AverageQuaternion(int numRotations,
   const Quaternion_t *sourceQ)
{
   Quaternion_t averageQ;
   int i;
```

/* initialise result with first rotation */

```
averageQ = (sourceQ[0]);
 /* loop over remaining rotations */
 for(i = 1; i<numRotations; ++i)
  /* Check for identical rotations and add to the average quaternion
  if(DotQuaternion(averageQ, sourceQ[i])>0.0f)
   averageQ = AddQuaternion (averageQ, sourceQ[i]);
  else
   averageQ = AddQuaternion (averageQ,
ReverseQuaternion(sourceQ[i]));
 /* Normalize */
 return NormalizeQuaternion (averageQ);
}
```

Does anyone have a suggestions on how do I can weight the different quaternions to get a weighted average rotation?

A previous post on averaging rotation matrices suggested:

- > I'd suggest transformation of the rotation matrices into
- > quaternions. The quaternion coefficients can be regarded as forming a
- > unit vector in 4-space. Your observations should give a cluster of
- > such vectors. The centroid of this cluster should give the mean
- > rotation.

I quite like this description, but I have no idea how to find the centroid of a cluster of vectors in 4-space. Can anyone point me to a book/paper/letter that a mere physicist might be able to understand.

Cheers. Graham