Subject: Averaging quaternions
Posted by GrahamWilsonCA on Thu, 18 Mar 2004 22:34:05 GMT
View Forum Message <> Reply to Message

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

Subject: Re: Averaging quaternions
Posted by jelansberry on Sat, 20 Mar 2004 15:23:03 GMT
View Forum Message <> Reply to Message

I've finally realized that all I contributed was questions and complaints and no alternative solutions.

If I were doing this, I would probably convert the quaternions to Euler (or Bryant) angles first (convert the quaternion to a direction cosine matrix, then extract the Euler angles). Then, I would compute the average of the Euler angles, and then convert the resulting average Euler angles back to a quaternion (convert the Euler angles to a direction cosine matrix, then extract the quaternion).

The only thing you have to worry about with Euler (or Bryant) angles is that there will always be a singularity for any chosen sequence. For example, if you choose a 3-2-1 Bryant sequence (i.e., first rotation about the 3-axis, second rotation about the subsequent 2-axis, third rotation about the subsequent 1-axis) then there will be a singularity whenever the second rotation angle is an odd multiple of 90 degrees (in that case, there is no unique solution for the first and third rotation angles). However, you can always look at the data and select an Euler (or Bryant) sequence that has no singularity.

John

"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

Subject: Re: Averaging quaternions
Posted by Arnold Neumaier on Sun, 21 Mar 2004 08:47:50 GMT
View Forum Message <> Reply to Message

jelansberry wrote:
> I've finally realized that all I contributed was questions and complaints
> and no alternative solutions.
>
> If I were doing this, I would probably convert the quaternions to Euler (or
> Bryant) angles first (convert the quaternion to a direction cosine matrix,
> then extract the Euler angles). Then, I would compute the average of the
> Euler angles, and then convert the resulting average Euler angles back to a
> quaternion (convert the Euler angles to a direction cosine matrix, then
> extract the quaternion).

This has exactly the same problems as averaging over quaternions, since angles are only unique up to a multiple of pi or 2pi; so the average depends on whether you represent an angle by a number close to pi or close to -pi ...

Arnold Neumaier

Subject: Re: Averaging quaternions
Posted by jelansberry on Sun, 21 Mar 2004 17:15:22 GMT
View Forum Message <> Reply to Message

"Arnold Neumaier" <Arnold.Neumaier@univie.ac.at> wrote in message
news:405D56B6.6030403@univie.ac.at...
> jelansberry wrote:
>> I've finally realized that all I contributed was questions and
complaints
>> and no alternative solutions.
>>
>> If I were doing this, I would probably convert the quaternions to Euler
(or
>> Bryant) angles first (convert the quaternion to a direction cosine
matrix,
>> then extract the Euler angles). Then, I would compute the average of
the
>> Euler angles, and then convert the resulting average Euler angles back
to a
>> quaternion (convert the Euler angles to a direction cosine matrix, then

>> extract the quaternion).
>
> This has exactly the same problems as averaging over quaternions, since
> angles are only unique up to a multiple of pi or 2pi; so the average
> depends on whether you represent an angle by a number close to pi or
> close to -pi ...
>
> Arnold Neumaier
>

"Uniqueness" of the Euler angles is not the issue, it's more an issue of continuity of the angles. Euler angles do not have the "same" problems as averaging over quaternions. My basic beef with averaging quaternions is that the initial result of the average is not a quaternion (i.e., the result does not have unit norm). Euler angles do not suffer from such a complication.

If all the OP is doing is trying to find the average attitude over some fairly small period of time, then one might expect the Euler angles corresponding to the quaternion samples to fairly continuous. Admittedly, if the quaternions are completely independent of one another, then such a continuity argument will fail. But then, what would be the purpose of finding an "average" attitude for quaternions that are randomly distributed?

I agree (and my post gave fair warning) that with Euler angles one has to be careful of choosing sequences near the singularity of the sequence. The problem you raise is essentially equivalent to that case - if you are near the singularity for the sequence, then you can expect large discontinuities in the extracted Euler angles. A quick plot of the Euler angles can help identify if the selected Euler (or Bryant) sequence is a "good" one. In general, it usually isn't that hard to avoid the singularity, particularly if you have an understanding of the underlying process that generated the quaternions in the first place.

John

Subject: Re: Averaging quaternions
Posted by Arnold Neumaier on Sun, 21 Mar 2004 17:48:33 GMT
View Forum Message <> Reply to Message

jelansberry wrote:
> "Arnold Neumaier" <Arnold.Neumaier@univie.ac.at> wrote in message
> news:405D56B6.6030403@univie.ac.at...
>
>> jelansberry wrote:
>>
>>> I would compute the average of the

>>> Euler angles, and then convert the resulting average Euler angles back to a
>>> quaternion (convert the Euler angles to a direction cosine matrix, then
>>> extract the quaternion).
>>
>> This has exactly the same problems as averaging over quaternions, since
>> angles are only unique up to a multiple of pi or 2pi; so the average
>> depends on whether you represent an angle by a number close to pi or
>> close to -pi ...
>>
>> Arnold Neumaier
>>
>
>
> "Uniqueness" of the Euler angles is not the issue, it's more an issue of
> continuity of the angles. Euler angles do not have the "same" problems as
> averaging over quaternions. My basic beef with averaging quaternions is
> that the initial result of the average is not a quaternion (i.e., the result
> does not have unit norm). Euler angles do not suffer from such a
> complication.

The real part of a unit quaternion (with nonnegative real part)
is redundant in that it can be recomputed from the imaginary part.
Thus averaging the imaginary parts and recomputing the real part
would be a simpler recipe of the same kind as yours with Euler angles.
And it would have exactly the same problems as the avarage-and-scale
method, although there are no asingularities. It is a matter of
non-uniqueness in both cases, which implies that one must make ad hoc
normalizations: A choice of sign in the quaternion case, and a choice
of some normalization interval in the Euler case. This cannot be
done without introducing discontinuities - these are not present
in the mathematics but only in the normalization chosen.

> If all the OP is doing is trying to find the average attitude over some
> fairly small period of time, then one might expect the Euler angles
> corresponding to the quaternion samples to fairly continuous.

Not if one of the angle is just a little less than pi and increasing
beyond pi (suddenly becoming -pi)

> I agree (and my post gave fair warning) that with Euler angles one has to be
> careful of choosing sequences near the singularity of the sequence.

AND near the normalization bounds! The average-and-scale technique is
thus even better since it has no singularities and only the
problem with possible discontinuities in the representation.

Arnold Neumaier

Subject: Re: Averaging quaternions
Posted by jelansberry on Sun, 21 Mar 2004 23:26:28 GMT
View Forum Message <> Reply to Message

"Arnold Neumaier" <Arnold.Neumaier@univie.ac.at> wrote in message
news:405DD571.1020208@univie.ac.at...
> jelansberry wrote:
>> "Arnold Neumaier" <Arnold.Neumaier@univie.ac.at> wrote in message
>> news:405D56B6.6030403@univie.ac.at...
>>
>>> jelansberry wrote:
>>>
>>>> I would compute the average of the
>>>> Euler angles, and then convert the resulting average Euler angles back
to a
>>>> quaternion (convert the Euler angles to a direction cosine matrix, then
>>>> extract the quaternion).
>>>
>>> This has exactly the same problems as averaging over quaternions, since
>>> angles are only unique up to a multiple of pi or 2pi; so the average
>>> depends on whether you represent an angle by a number close to pi or
>>> close to -pi ...
>>>
>>> Arnold Neumaier
>>>
>>
>>
>> "Uniqueness" of the Euler angles is not the issue, it's more an issue of
>> continuity of the angles. Euler angles do not have the "same" problems
as
>> averaging over quaternions. My basic beef with averaging quaternions is
>> that the initial result of the average is not a quaternion (i.e., the
result
>> does not have unit norm). Euler angles do not suffer from such a
>> complication.
>
> The real part of a unit quaternion (with nonnegative real part)
> is redundant in that it can be recomputed from the imaginary part.
> Thus averaging the imaginary parts and recomputing the real part
> would be a simpler recipe of the same kind as yours with Euler angles.
> And it would have exactly the same problems as the avarage-and-scale
> method, although there are no asingularities. It is a matter of
> non-uniqueness in both cases, which implies that one must make ad hoc
> normalizations: A choice of sign in the quaternion case, and a choice
> of some normalization interval in the Euler case. This cannot be
> done without introducing discontinuities - these are not present
> in the mathematics but only in the normalization chosen.

I'm not following what you mean by a "normalization" interval (or "normalization bounds" below). Euler angles do not require "normalization." I still maintain that any "discontinuity" in the Euler angles usually is the result of passing through the singularity that exists in any three-parameter representation of the direction cosine matrix, and it is generally easy to avoid this situation. Let's put it this way - in the aerospace industry (in which I work), we typically use the 3-2-1 Bryant sequence for guided missiles. Why? Because the singularity (i.e., where yaw and roll are not distinguishable) occurs at a pitch angle of 90 degrees, which is an unusual attitude for a guided missile to obtain (relative to local level). One could say the same for a manned aircraft. You don't see too many planes heading "straight up" relative to local level. If one stays away from 90 degree pitch angles, a plot of the Bryant angles over time is rarely discontinuous. Hence, averaging the Euler angles over any time interval is usually a reasonable thing to do. For that matter, interpolating in a time history of Euler angles is usually a reasonable thing to do. It is rarely a problem, and, if it is, selecting an alternative Euler/Bryant sequence is typically enough to avoid any problems.

As for quaternions, even if cos(theta/2) passes through zero, it does not represent a singularity or discontinuity unless one desires to keep cos(theta/2) positive (always). I've yet to find a simulation where the time history of quaternions shows any singularity or discontinuity. That's one of the advantages of quaternions relative to Euler angles when it comes to simulations.

>
>> If all the OP is doing is trying to find the average attitude over some
>> fairly small period of time, then one might expect the Euler angles
>> corresponding to the quaternion samples to fairly continuous.
>
> Not if one of the angle is just a little less than pi and increasing
> beyond pi (suddenly becoming -pi)

Yes, typically this happens when you pass through the singularity of the Euler sequence. It rarely occurs otherwise. This requires inspection of the data and a "wise" choice of Euler sequence, but what's the big deal? If your main argument is that one of the Euler angles could pass from -pi to pi (e.g., the yaw or roll Euler angle getting close to +/- 180 degrees), then the argument is weak. I can always redefine my reference frame to avoid that kind of problem for the purpose of "averaging" Euler angles. It really isn't that hard - I do it all the time.

>
>> I agree (and my post gave fair warning) that with Euler angles one has
to be
>> careful of choosing sequences near the singularity of the sequence.
>
> AND near the normalization bounds! The average-and-scale technique is

> thus even better since it has no singularities and only the
> problem with possible discontinuities in the representation.

As I said previously, my main issue with "average and scale" is that the initial "average" results in something that is not even a quaternion. This is why SLERP was created in the first place - to interpolate between two unit vectors under the constraint that the result is also a unit vector. If the unit vectors are widely separated (the angle between them is "large"), linear interpolation between the unit vector components (followed by "normalization" to get a unit vector back) may result in something non-sensical. I can imagine the same type of problems with "averaging and scaling" a bunch of quaternions. Clearly, this will be less of a "problem" if the unit vectors are not widely separated, which is probably the case for the OP. Admittedly, if the unit vectors are close to one another, it probably doesn't matter much what you do. But I think I can take the average of Euler angles of 0 degrees and 135 degrees and get a more reasonable result than averaging the components of the corresponding quaternions, especially if the data is taken at regularly sampled time intervals.

>
>
> Arnold Neumaier
>

Subject: Re: Averaging quaternions
Posted by Arnold Neumaier on Mon, 22 Mar 2004 08:54:59 GMT
View Forum Message <> Reply to Message

jelansberry wrote:

> I'm not following what you mean by a "normalization" interval (or
> "normalization bounds" below). Euler angles do not require "normalization."
> I still maintain that any "discontinuity" in the Euler angles usually is the
> result of passing through the singularity that exists in any three-parameter
> representation of the direction cosine matrix, and it is generally easy to
> avoid this situation.

No. One has this phenomenon alreasy in 2D rotations. The rotation angle is determined only up to a multiple of 2 pi, and one has to normalize the angle by, say, forcing it into the interval]-pi,pi]. Then you get problems in averaging two very close rotation angles, one just below pi, one just abouve -pi. The result will be close to zero instead of close to +-pi.

If you choose as normalization interval [0,2 pi[, the same problem happens when averaging a tiny rotation to the right (angle=eps) and to the left (angle=2 pi-eps).

There is no way to avoid such discontinuities, and one has the same
problems in 3 dimensions, and with quaternions.

But the quaternion recipe I gave a few days ago works perfectly
if the rotations to be averaged are not too much scattered.

Arnold Neumaier

Subject: Re: Averaging quaternions
Posted by Arnold Neumaier on Tue, 30 Mar 2004 10:53:06 GMT
View Forum Message <> Reply to Message

Graham wrote:
> 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.
>
>> 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:

Yes, that works indeed: it suffices to reorient the quaternions
that have negative inner product with one of them.
This is a nice observation that simplifies the above and makes
it faster.

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

In the above, simply weight your quaternions before adding them.
Thus the final algorithm is:

1. orient all quaternions to positive inner product with the first one
2. sum the results, weighted by their importance
3. normalize to unit norm

4. If some of the inner products in step 1 were in [-0.5,0.5], it is possible that some orientation went wrong. In this case, one should repeat the cycle with the result of step 3 in place of 'the first one', and skip in this second round in step 2 all quaternions with an inner product in [-0.5,0.5] as too scattered.

If you also want to get an assessment of the accuracy of the final result, more care is needed. (Projection to the tangent plane etc.)

> 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.

This is essentially the same as above, except that the reorientation is missing (which is essential, however).

Arnold Neumaier