Getting Started with the Unscented Kalman Filter (UKF) in ODTK

1 Introduction

You can configure the Filter process in ODTK to operate as an Extended Kalman Filter (EKF), an Unscented Kalman Filter (UKF), or as a hybrid between the EKF and UKF. The EKF is the heritage filtering capability, while the UKF and Hybrid options were introduced in ODTK 7.0. There are some limitations to the current capabilities, however, that you should understand in order to use the tool properly.

You can consider a generic sequential filter process to be a recursion of time and measurement updates. Time updates serve to transition the state and state uncertainty estimates from one time to another time, typically forward in time in the filter process and backward in time in the smoother process. Measurement updates use observational information at a point in time to produce updated state and state uncertainty estimates that incorporate the knowledge contained in the observations.

The EKF uses a single nonlinear propagation of the state and a linear mapping of the state-error covariance to transition through time. During the measurement update, the EKF uses a single nonlinear measurement model evaluation and a linear mapping of the measurement residual to update the state and state uncertainty estimates.

The UKF uses two or more nonlinear propagations of the state to transition both the mean estimate and state-error covariance through time[1]. During the measurement update, the UKF uses two or more nonlinear measurement model evaluations to compute the expected mean residual and residual error variance followed by a linear update of the state and state uncertainty estimates.

2 Filter configuration

The EKF is the default filter configuration in ODTK. You can specify other configurations, including the UKF and Hybrid formulations, using the ProcessControl.HigherOrderCorrections settings in the Filter properties. To specify the UKF, set the MitigationMethod to Unscented Filter. With the MitigationMethod set, you can specify the full UKF by setting the UnscentedFilter.UpdateTypes setting to MeasurementAndTime. This setting, along with specification of values for Alpha, Beta, and Kappa, dictate that the scaled, unscented transform will be used in both the time and measurement updates of the filter. Alternatively, there are some cases where the unscented transform may be desirable in the measurement update, but does not provide any advantage in the time update. The use of satellite-to-satellite observations inside a formation where absolute orbit uncertainty is not large is one such example. Under these conditions, setting UnscentedFilter.UpdateTypes to Measurement can save processing time[2].

3 Unscented transform

ODTK uses the scaled, unscented transform in the UKF formulation of the time and measurement updates. Generally speaking, the unscented transform uses nonlinear transformations of a set sample points to convert a mean and covariance in one set of coordinates to a mean and covariance in a different set of coordinates. In the time update, the transformation consists of the time transition of the estimation state, where the state elements at time t1 represent the initial coordinates and the state elements at time t2 represent the final coordinates. In the measurement update, the transformation moves between state and measurement coordinates. In both cases, ODTK performs (2N+1) nonlinear transformations, where N is the size of the estimation state.

The scaled, unscented transform utilizes sample points that ODTK constructs based on the mean and covariance of an initial distribution. These sample points are dependent upon the state size, N, and the specification of two parameters that serve to scale the offset between the sample (sigma) points and the mean.

- α – This is a scale factor for the offsets between sigma points from the unscaled scented transform and the mean. Since offsets grow with the size of the estimation state, it is important to control the size of the offsets so that local behavior is properly captured in the transformation. Choosing alpha as proportional to $1/\sqrt{N}$ is an effective strategy.
- κ – This is a secondary scale factor in the unscented transform, typically set to zero. Use of negative values can lead to nonpositive definite error covariance.

4 Estimation

Use the UKF and Hybrid for estimation in the same manner as the EKF in ODTK.

4.1 Process noise

The method for application of process noise in the EKF involves the division of time update intervals into subintervals based on the ProcessNoiseUpdateInterval setting in the Filter properties. The Filter adds accumulated process noise to the state-error covariance at the end of each subinterval. This allows the process noise to feed into the propagation of the overall covariance in a manner that is temporally and geometrically appropriate.

In the UKF, using the process noise strategy describe above can be problematic. Addition of process noise into the covariance requires that the UKF do the following:

- Reconstruct the covariance from the sample points
- Add in the process noise
- Resample the resulting covariance to yield new sample points

In certain cases, such as in the estimation of highly eccentric orbits, this procedure can result in much larger than expected uncertainty growth[2]. To overcome this problem while still constructing process noise under appropriate conditions, ODTK provides the UnscentedFilter.Resample-

OnProcessNoiseGrid option in the Filter properties. When you disable this option, the Filter accumulates the process noise parallel to the propagation of the trajectory across the overall time update interval and only adds it to the covariance when needed to process a measurement.

4.2 Smoothing

The UKF generates *.rough files in support of smoothing. However, ODTK smoothing operations do not use the unscented transform during the smoothing processing. Results obtained using a linear smoother on the output from the UKF may, therefore, be suboptimal.

4.3 Restart

You can use the UKF in restart mode. In this case, the UKF reads in the initial estimation state and state-error covariance from the selected restart record. If the selected restart record is on a time where the UKF processed an observation, as would be the case if you set the stop time to be the time of the last observation, the subsequent estimation run should produce results identical with a run that covered both the prior and current time spans. If, however, the selected restart record is at a time between observations, and therefore would have been inside a time update on a continuous estimation run, then the results of the second filter run will be slightly different than a run which covered both the prior and current time spans. This difference comes from the need to resample the covariance at the start of the second filter run.

4.4 State differencing

You can use the state differencing function (StateDiffTool.htm) to difference run files generated while using the UKF capability.

---

[1] Julier, S.J., Uhlmann, J.K. and Durrant-Whyte, "A new approach for filtering nonlinear systems," 1995 American Control Conference, Seattle, WA, pp. 1628-1632.

[2] Woodburn, J. and Coppola, V., "Analysis of Relative Merits of Unscented and Extended Kalman Filters in Orbit Determination," AAS 19-781, 2019 AAS/AIAA Astrodynamics Specialists Conference, Portland, ME, August 2019.