Stone Soup: an open source tracking and state estimation framework; principles, use and applications

ISIF Fusion 2023 tutorial proposal

Duration:

Half-day

Presenters:

Lyudmil Vladimirov (University of Liverpool, UK)

Lyudmil is a Postdoctoral Research Associate at the University of Liverpool. His research focuses on Bayesian target tracking and intent prediction algorithms, with specific interest on the applications of Sequential Monte Carlo algorithms (i.e., particle filter, smc samplers) in the maritime domain.

Peter Carniglia (Defence Research and Development Canada, Canada)

Peter works at DRDC Ottawa as a Defence Scientist. His area of specialization is in drone detection and tracking, sensor data fusion and Bayesian state estimation.

James Wright (Defence Science and Technology Laboratory, UK)

James is a senior data fusion scientist focussing on developing the algorithmic offerings within the Stone soup framework and applying the framework through simulation modelling and analysis of real-world data in the maritime and air domains.

Nicola Perree (Defence Science and Technology Laboratory, UK)

Nikki is a data scientist at Dstl. Her work involves developing the Stone Soup framework with a particular focus on autonomous sensor management capabilities.

Henry Pritchett (Defence Science and Technology Laboratory, UK)

Henry is a Graduate Scientist at Dstl. His work has focussed on the creation of real world simulation using Stone Soup, as well as the analysis of state estimation algorithms for the PNT use-case.

Website:

https://stonesoup.readthedocs.io/

Target audience:

The tutorial will be suitable for students new to statistical inference and estimation; recent graduates in the mathematical sciences moving into tracking and state estimation; practitioners in industry and government with an interest in algorithm comparison and applications; academic researchers for whom robust baselining is necessary to demonstrate the efficacy of their work.

Abstract:

It is currently difficult and time consuming for academic researchers to recreate state-of-the-art tracking and state estimation algorithms to benchmark their work. Comparison of new algorithms with existing solutions involves recoding algorithms from the literature. Industrial users also find it difficult to assess which algorithms meet their, often quite varied, requirements.

The Stone Soup framework is designed to provide a flexible and unified software platform for researchers and engineers to develop, test and benchmark a variety of existing multi-sensor and multi-object estimation algorithms. It is also designed to allow rapid prototyping of new algorithms in high-level languages (e.g. Python, Matlab), as well as development in compiled languages (e.g. C++), by providing a

set of libraries which implement the necessary functions for tracking and state estimation. It profits from the object-oriented principles of abstraction, encapsulation and modularity, allowing users (beginners, practitioners or experts) to focus only on the most critical aspects of their problem.

Stone Soup is endorsed by ISIF's working group on Open Source Tracking and Estimation (OSTEWG). These tutorials will introduce participants to Stone Soup's basic components and how they fit together. They are delivered by way of demonstrations, set tasks and interactive tutorials where participants will be encouraged to write and modify algorithms. These tasks will be written up in the form of interactive browser-based applications which combine the ability to run code with a presentation environment suitable for instruction.

The tutorial will begin with basic examples using linear transition models, abstract range-bearing sensors and single-targets using the extended and unscented Kalman and particle filters. Multiple targets, clutter and methods of data association will be introduced. The final part of the tutorial will look at applying Stone Soup to several scenarios involving simulated and real-world data.

Course outline:

Introduction to Stone Soup

This session will start with an introduction to Stone Soup, before covering an introduction to single and multi-target tracking, and tracking practicalities.

1. Single Target Tracking

This lesson starts by introducing some of the basic features of Stone Soup using a single sensor, single target scenario. It starts with the linear-Gaussian (Kalman) filter and proceeds to introduce the most commonly-used non-linear alternative, the extended Kalman filter (EKF). In the final part of this lesson we look at a class of sequential Monte Carlo sampling methods, and in particular, the particle filter.

2. Multi Target Tracking

In this lesson, we introduce the use of data association algorithms in Stone Soup and demonstrate how they can mitigate confusion due to clutter, and also enable multi-target tracking.

To begin with we use nearest neighbour methods, which are conceptually simple and make firm associations between predictions and detections. Making a firm assignment at each time-step between a single track and a single measurement can be problematic. So we will also look at a probabilistic measure of assignment and see how such methods increase the robustness of the multi-target tracking solution.

3. Tracking Practicalities

In practical tracking applications, targets may appear and disappear. This could be because they enter/exit the sensor field of view, are occluded or obscured, or become undetectable. It is of key importance to understand how well an algorithm copes with these phenomena. This lesson looks at how to cope with the challenges of non-observable targets before introducing some of the available metrics in Stone Soup and how they can be used to assess multi-target tracking performance.

4. Applications of Stone Soup

In this final section, we will demonstrate and apply Stone Soup with real data sets, with practical examples. This could include video processing and tracking; drone tracking with radar; using real flights path data (ADS-B) with simulated sensors; sensor management in simulated environment.

Requirements:

Laptop
Download, install, test: Python 3.8+, Jupyter, Stone Soup (step by step instructions will be provided in advance)
Notebooks will be distributed in advance.