

Implementation of AKKF-based Multi-Sensor Fusion Methods in Stone Soup

James S. Wright, Mengwei Sun, Mike E. Davies, Ian K. Proudler, James R. Hopgood

Abstract—This paper explores the increasing demand for accurate and resilient multi-sensor fusion techniques, particularly within 3D tracking systems enhanced by drone technology. Employing the adaptive kernel Kalman filter (AKKF) methodology within the Stone Soup framework, our research seeks to develop robust fusion approaches capable of seamlessly amalgamating data from a multi-sensor arrangement with fixed ground sensors and dynamic sensors mounted on drones. By capitalising on the adaptive nature of the AKKF, we aim to refine the precision and dependability of 3D object tracking in intricate scenarios. Through empirical evaluations, we illustrate the effectiveness of our proposed AKKF-based fusion strategies in enhancing tracking performance within the Stone Soup framework, thus contributing to the advancement of multi-sensor fusion methodologies within this framework.

Index Terms—Adaptive kernel Kalman filter; sensor fusion, Stone Soup; 3D Tracking.

I. INTRODUCTION

In recent years, there has been a notable increase in demand for accurate and robust multi-sensor fusion techniques, particularly within 3D tracking systems employing drones [1]. This surge can be attributed to the proliferation of sensor technologies and their integration into various applications, from autonomous vehicles to surveillance systems. Drones, equipped with sensors, offer distinct advantages in tracking scenarios, providing dynamic viewpoints, swift repositioning capabilities, and extensive coverage across large areas [2], [3]. By integrating sensors on drones, a complementary relationship is established with fixed ground sensors, extending

tracking range, enhancing visibility, and enabling adaptive tracking strategies in dynamic environments. Furthermore, drones possess the flexibility to navigate obstacles and access remote or inaccessible areas, thereby improving the overall effectiveness of the tracking system.

This paper addresses multi-sensor tracking scenarios involving fixed sensors on the ground and dynamic sensors mounted on drones. This hybrid setup introduces unique challenges, including disparate sensor characteristics, varying observation geometries, and dynamic environmental conditions. Among the plethora of fusion methods, the adaptive kernel Kalman filter (AKKF) has emerged as a promising approach with favourable computational complexity and offers an excellent opportunity for parallelization [4], [5]. Herein, we concentrate on implementing AKKF-based fusion methods within the framework of Stone Soup [2], [6]. Our work builds upon the foundations laid by the seminal paper “Adaptive Kernel Kalman Filter Multi-Sensor Fusion” [7], which introduced the theoretical underpinnings of AKKF-based fusion techniques.

We implement the published algorithms within the Stone Soup platform for several reasons. Stone Soup offers a versatile framework for algorithm development and evaluation, ensuring reproducibility and transparency. The modular architecture of Stone Soup allows for comprehensive testing across different scenarios. Integrating the proposed algorithm into Stone Soup extends its accessibility and facilitates comparative analyses with other state-of-the-art methods.

Through empirical evaluations and comparative analyses, we demonstrate the efficacy of the proposed AKKF-based fusion methods in enhancing tracking performance across various scenarios, leveraging the contributions of flying drones. Our contributions advance the state-of-the-art in multi-sensor fusion techniques and provide valuable insights into the practical implementation of AKKF-based fusion methods within the context of the Stone Soup framework.

In the subsequent sections, we delve into the technical details of our approach, starting with a brief overview of the AKKF-based fusion algorithms. We then describe the experimental setup, present our results, and discuss the implications of our findings. Finally, we conclude with remarks on future research directions and potential applications of our work.

II. APPLICATION OF THE AKKF-BASED FUSION METHODS IN TARGET TRACKING

In the depicted tracking scenario, a single target manoeuvres within the environment, as shown in Fig. 1, which is generated using the Stone Soup framework. Alongside fixed

J. S. Wright is with Dstl, Dstl Porton Down, Bldg 5, Rm 101, Salisbury, Wiltshire, SP4 0JQ, U.K. E-mail: jwright2@dstl.gov.uk.

M. W. Sun is with the School of Aerospace, Transport and Manufacturing (SATM), Cranfield University, College Road, Cranfield, MK43 0AL, U.K. E-mail: mengweisun@cranfield.ac.uk.

M. E. Davies and J. R. Hopgood are with Institute of Imaging, Data, and Communications, University of Edinburgh, Edinburgh, EH9 3FG, U.K. E-mail: (mike.davies; james.hopgood)@ed.ac.uk.

I. Proudler is with the Centre for Signal & Image Processing (CeSIP), Department of Electronic & Electrical Engineering, University of Strathclyde, Glasgow, G1 1XW, U.K. E-mail: ian.proudler@strath.ac.uk.

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) Grant number EP/S000631/1; and the MOD University Defence Research Collaboration (UDRC) in Signal Processing.

The contents include material subject to ©Crown copyright (2024), Dstl. This material is licensed under the terms of the Open Government Licence except where otherwise stated. To view this licence, visit <http://www.nationalarchives.gov.uk/doc/open-government-licence/version/3> or write to the Information Policy Team, The National Archives, Kew, London TW9 4DU, or email: psi@nationalarchives.gov.uk

For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.




Fig. 1: Integrated tracking system: single target tracking with fixed ground sensor and drone in Stone Soup.

roadside base stations with predetermined positions, drones are deployed to enhance monitoring and tracking. These drones are equipped with global navigation satellite system (GNSS) to precisely pinpoint their locations within the environment and enable them to measure their altitude. As the target moves, drones take flight to oversee and track the target's movements using different sensors, such as bearing-only or range-bearing sensors. In this paper, fixed-wing drones flying in a circle of constant radius are employed for tracking due to their prolonged endurance, higher speeds and stability [8]. This integration of drones introduces a dynamic element to the tracking system, enabling flexible and adaptable surveillance as the target traverses the area. By combining data from fixed base stations and drones, the tracking system achieves comprehensive coverage and real-time accuracy in monitoring the target's trajectory and movements.

Consider the following generic (linear or non-linear) dynamical state space model (DSSM), which can mathematically describe this tracking scenario:

$$\mathbf{x}_k = f(\mathbf{x}_{k-1}, \mathbf{u}_k), \quad k = 1, \dots, K \quad (1)$$

$$\mathbf{y}_{k,n} = h_{k,n}(\mathbf{x}_k, \mathbf{v}_{k,n}), \quad n = 1, \dots, N. \quad (2)$$

Here, \mathbf{x}_k represents the hidden state of the target at time slot k , $\mathbf{y}_{k,n}$ represents the observed signal from the target at the n -th sensor node. The motion function is denoted as $f(\cdot)$ and the measurement function of the n -th sensor node is denoted as $h_n(\cdot)$, with \mathbf{u}_k and $\mathbf{v}_{k,n}$ representing the process and measurement noise, respectively. The total number of sensors $N = N_D + N_G$, N_D represents the count of dynamic sensors equipped on drones, while N_G represents the number of fixed ground sensors.

For model-based target tracking/localisation, Bayesian filter-based tracking and triangulation-based localisation are commonly used methods. Bayesian filter-based tracking utilises probabilistic models and recursive estimation, making it well-suited for dynamic tracking and continuous updates.

Triangulation-based localisation uses geometric principles and direct calculations from known references, typically applied for static position estimation or with static reference points. Bayesian filters can work with varying numbers of sensors and excel in sensor fusion, whereas triangulation-based localisation requires a specific number of sensors to solve the geometric equations necessary for determining the object's position; the position cannot be uniquely determined if fewer sensors are available than the minimum required. Based on these differences, we estimate the hidden state recursively from a Bayesian perspective to address the tracking scenario depicted in Fig. 1. Specifically, the estimation of hidden states \mathbf{x}_k is solved by constructing the posterior probability density function (PDF) of hidden states based on all available information, including the DSSM and received measurements, i.e., updating $p(\mathbf{x}_k | \mathbf{y}_{1:N,1:k})$. The AKKF uses the reproducing kernel Hilbert space (RKHS) to transform the received measurements into RKHS to address the issues of geometry arising from triangulation based methods. In paper [7], we proposed two fusion methods, which are summarised as follows.

A. Centralised Fusion

In centralised fusion, all current measurements $\mathbf{y}_{k,1:N}$ and measurement models $h_{k,1:N}(\cdot)$ are transmitted from sensor nodes to the fusion center (FC). The global measurement model is then formulated at the FC as

$$\mathbf{Y}_k = H_k(\mathbf{x}_k, \mathbf{V}_k) \quad (3)$$

where $\mathbf{Y}_k = [\mathbf{y}_{k,1}, \dots, \mathbf{y}_{k,N}]^T$. The centralised fusion is realised by applying the AKKF directly at the FC [7].

B. Semi-decentralised Fusion

Unlike centralised fusion, where all processing occurs exclusively at the FC and requires real-time transmission of data from all sensors to the FC, the semi-decentralised approach proposed in [7] distributes local processing tasks among individual sensors. Here, each sensor computes local posterior kernel weight vectors and covariance matrices based on individual measurements. The global kernel mean embedding (KME) is then approximated at the FC using a weighted combination of these local estimates, following Eqn. (37) and Eqn. (38) in [7]. This semi-decentralised scheme not only alleviates the computational load at the FC but also mitigates bandwidth requirements by eliminating the need for real-time transmission of all data to the FC. Additionally, it offers a more modular approach to sensor networks, as the FC does not require detailed knowledge of individual sensor characteristics.

III. IMPLEMENTATION IN STONE SOUP

Stone Soup [6], [9] is an open-source tracking and state estimation framework following an object-orientated modular approach to build tracking algorithms. This allows the utilisation of algorithmic components in a plug-and-play manner without requiring a full understanding of the algorithms [10]. In this section, we will introduce the flow diagrams that implement AKKF-based centralised and semi-decentralised fusion

within Stone Soup, emphasizing their unique characteristics and differences from the standard AKKF. Subsequently, we will elucidate the new components devised in Stone Soup to create these fusion schemes.

A. Flow Diagrams for AKKF Fusion in Stone Soup

The standard AKKF for single sensor single target tracking, depicted in Fig. 5(a), operates across state space, measurement space, and kernel space, encompassing three key modules. Firstly, a predictor integrates prior and proposal information at time $k - 1$. Subsequently, an updater employs the predicted values to refine the kernel weight and covariance. Finally, another updater generates the proposal state particles.

The centralised fusion AKKF, as illustrated in Fig. 5(b), is a minor departure from the standard AKKF due to its centralised processing approach. In this configuration, all three modules, predictor, updater, and proposal particle generation, are consolidated at the FC. This centralised architecture enables information aggregation from multiple sensors, offering advantages regarding global perception and decision-making. One key distinction from the standard AKKF is the handling of input measurements. Instead of processing measurements from individual sensors, the centralised fusion AKKF operates on a combined vector comprising measurements from all sensors. Moreover, the measurement model employed in the centralised fusion AKKF is a global representation that integrates information from all sensor measurement models. By incorporating measurements and measurement model information into a unified framework, the centralised fusion AKKF can comprehensively view the target's state and effectively account for variations in sensor characteristics and environmental conditions, resulting in more reliable estimation outcomes.

As depicted in Fig. 5(c), the semi-decentralised fusion AKKF operates at both local sensors and the FC. Initially, the proposal particle generation and prediction modules are processed at the FC. Subsequently, the update step across measurement and kernel space is processed locally at the sensor, culminating in calculating the updated kernel weight mean vector and covariance matrix based on the locally received measurements. These locally updated parameters are then transmitted back to the FC for fusion, employing the weighted Kullback-Leibler average (KLA) mechanism. The semi-decentralised fusion AKKF introduces a hybrid approach that combines decentralised processing at local sensors with centralised fusion at the FC. Unlike the standard and centralised fusion AKKFs, which operate solely centrally, this approach enables local sensor updates, reducing computational burden and enhancing real-time responsiveness. Locally updated parameters are transmitted to the FC for centralised fusion using the weighted KLA, ensuring the final estimate reflects the most reliable information from all sensors.

B. New Components in Stone Soup

Previous work [2] established the groundwork for implementing standard AKKF within a single sensor context. The components integrated the capabilities of

kernels, kernel states, the AKKF predictor and the AKKF updater. Particularly of use in the context of the fusion methods is the `AdaptiveKernelKalmanUpdater`. This paper extends prior efforts by introducing the *Updaters* required for multi-sensor fusion: the `CentralisedAdaptiveKernelKalmanUpdater` and the `SemiDecentralisedAdaptiveKernelKalmanUpdater`. These two new components allow for the centralised and semi-decentralised fusion methods to take measurements from multiple sensors, process the measurements at the sensor and communicate either the measurements for the centralised case or the updated kernel weights and covariance to the fusion centre in the fusion tracking loop, as depicted in the flow diagram in Fig. 5.

The `CentralisedAdaptiveKernelKalmanUpdater` inherits from the `AdaptiveKernelKalmanUpdater` and modifies the `update()` method to include the set of hypotheses (the prediction, measurement pairs) for each sensor. The update method continues with the combined sensor measurement information following the standard AKKF Update. This is the only required change to implement the centralised fusion method.

The `SemiDecentralisedAdaptiveKernelKalmanUpdater` similarly inherits from the `AdaptiveKernelKalmanUpdater` and modifies the `update()` method. The Semi-decentralised fusion processes the measurement on the sensor before sending the mean vector and covariance matrix of kernel weights information to the fusion centre. This is implemented by running parallel single sensor *Updaters* before fusing the combined kernel weights' mean and covariance at the FC.


IV. SIMULATION RESULTS

This paper uses constant-velocity (CV) motion model, following the `ConstantVelocity` class as shown in [11], to approximate the target trajectory. For sensor configuration, bearing-only sensors are deployed to gather observations from the target. Drones are equipped with 3D angle-only sensors, such as omnidirectional cameras capable of providing monocular vision. These cameras capture the target's movements from various angles. On the ground, 2D bearing-only sensors, such as angle of arrival sensors, are utilised to obtain the arrival angle from the target's perspective. The bearing-only observation models for the drones and ground sensors are defined as (4a) and (4b), respectively [12].


$$\mathbf{y}_{k,n}^D = \begin{pmatrix} \theta_{k,n}^D \\ \phi_{k,n}^D \end{pmatrix} = \begin{pmatrix} \arctan \frac{\eta_k - \eta_{k,n}^D}{\xi_k - \xi_{k,n}^D} \\ -\arctan \frac{\zeta_k - \zeta_{k,n}^D}{\sqrt{(\eta_k - \eta_{k,n}^D)^2 + (\xi_k - \xi_{k,n}^D)^2}} \end{pmatrix} + \mathbf{v}_{k,n}^D, \quad (4a)$$

$$y_{k,n}^G = \theta_{k,n}^G = \arctan \frac{\eta_k - \eta_n^G}{\xi_k - \xi_n^G} + v_{k,n}^G. \quad (4b)$$


Here, $\xi_{k,n}^D$, $\eta_{k,n}^D$ and $\zeta_{k,n}^D$ represent the coordinates of the n -th drone sensor, where $n = 1, \dots, N^D$, ξ_n^G , η_n^G represent the coordinates of the ground sensor, where $n = 1, \dots, N^G$. The drones are equipped with GNSS, providing real-time information for $\xi_{k,n}^D$, $\eta_{k,n}^D$ and $\zeta_{k,n}^D$. The 3D bearing $\mathbf{y}_{k,n}^D$ consists




(a)



(b)




(c)




(d)


Fig. 2: Tracking performance using different methodologies: (a) Centralised fusion; (b) Semi-decentralised fusion; (c) The drone sensor only; (d) The ground sensor only.




(a)



(b)



(c)



(d)

Fig. 3: Tracking performance using different methodologies in Stone Soup: (a) Centralised fusion; (b) Semi-decentralised fusion; (c) The drone sensor only; (d) The ground sensor only.

of two components: $\theta_{k,n}^D$ and $\phi_{k,n}^D$ provide the vertical angle and horizontal angle relative to the drone's position, respectively. The 2D bearing observation $y_{k,n}^G$ provides the angle relative to the sensor's position on the ground. ξ_k , η_k and ζ_k represent the current coordinates of the target. As the target is assumed to be moving on the flat ground, $\zeta_k = 0$. The measurement noise $\mathbf{v}_{k,n}^D$ for drone sensors and $v_{k,n}^G$ for ground sensors, are assumed to follow Gaussian distributions with zero mean, $\mathbf{v}_{k,n}^D \sim \mathcal{N}(\mathbf{0}, \sigma_D^2 \mathbf{I}_2)$, $v_{k,n}^G \sim \mathcal{N}(0, \sigma_G^2)$. These noise terms capture the random fluctuations or errors introduced during measurement.

In this simulation, the timestep in the motion model is $\Delta T = 1s$, and the magnitude of the noise is set to be $1e-2$. We consider a scenario featuring one ground sensor and one drone sensor, with the measurement noise variance set as $\sigma_G^2 = 2.6e-4$ and $\sigma_D^2 = 6.1e-3$. To initialise the tracking, the prior distribution of the hidden state, i.e., \mathbf{x}_0 , is set to be Gaussian with mean $\bar{\mathbf{x}}_0 = [-0.05, 0.001, 0.7, -0.05]^T$ and covariance matrix $\mathbf{P}_0 = \text{diag}[1e-2, 2.5e-5, 1e-2, 1e-4]$. Figures 2 and 3 display the tracking performance obtained by centralised fusion, semi-decentralised fusion, the drone sensor only, and the ground sensor only using the quadratic kernel-based AKKFs with the particle size set to 100, in MATLAB 2023a and Stone Soup respectively. Additionally, in Fig. 4, we compare the average logarithmic mean squared error (LMSE) performance obtained by running 100 Monte Carlo (MC) simulations with particle sizes of 100 and 200.

From these simulation results, it is evident that filters using a single sensor measurement lose track. The performance of the ground sensor is inferior, mainly due to the sensor's fixed position at the end of the trajectory, which constrains the sensor's perspective and results in insufficient position information based solely on received bearing measurements. The drone, with its ability to cover a larger surveillance area and exhibit higher flexibility than the fixed ground sensor, offers distinct advantages in tracking scenarios. This is reflected in the enhanced tracking performance observed when utilising drone-based measurements. Fusion-based filters demonstrate improved performance, highlighting the benefit of data fusion. Further, the results exhibit a remarkable consistency between the MATLAB 2023a and Stone Soup frameworks, demonstrating the validity and effectiveness of Stone Soup as a versatile tool for addressing complex tracking challenges across various scenarios and environments. Additionally, the Stone Soup framework's ability to replicate results from MATLAB simulations further validates its suitability for practical deployment in real-world tracking and surveillance applications.

V. CONCLUSIONS

This paper provides an implementation of fusion methods of the AKKF algorithm in Stone Soup. We utilised and extended existing AKKF components to incorporate the fusion updaters. The new fusion algorithms provide extensions to the existing AKKF algorithm for the multi-sensor use cases. Additionally, this paper demonstrated the use of the AKKF algorithm with a mixture of sensor modalities including moving sensors. This




Fig. 4: Average LMSE performance with different particle sizes. The error bars are calculated as $E(\text{LMSE}) \pm \text{Std}(\text{LMSE})$.

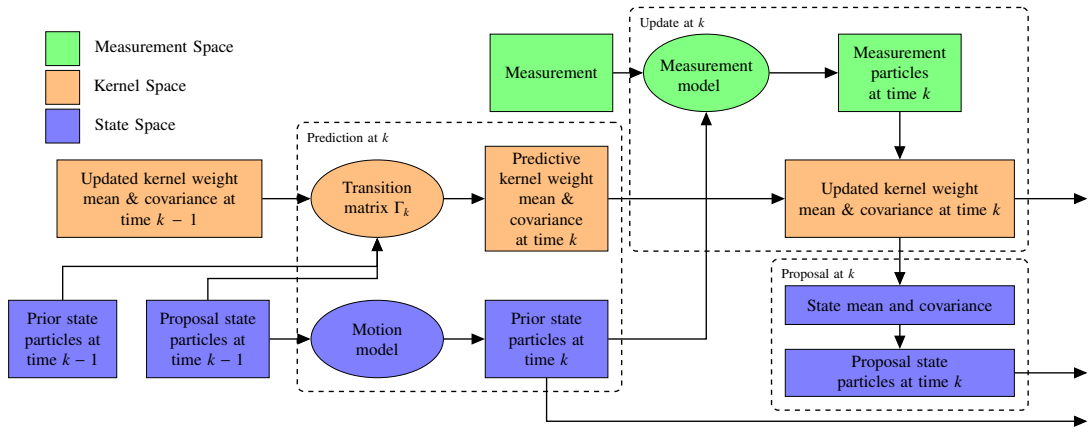
showed that Stone Soup is a very powerful framework for implementing sophisticated scenarios, i.e., implementing 2D and 3D scenarios, and decreases the complexity of introducing additional sensors and sensor modalities.

The authors are interested in extending the AKKF further [5], [13] by offering additional examples on the AKKF fusion methods in the Stone Soup documentation [9].

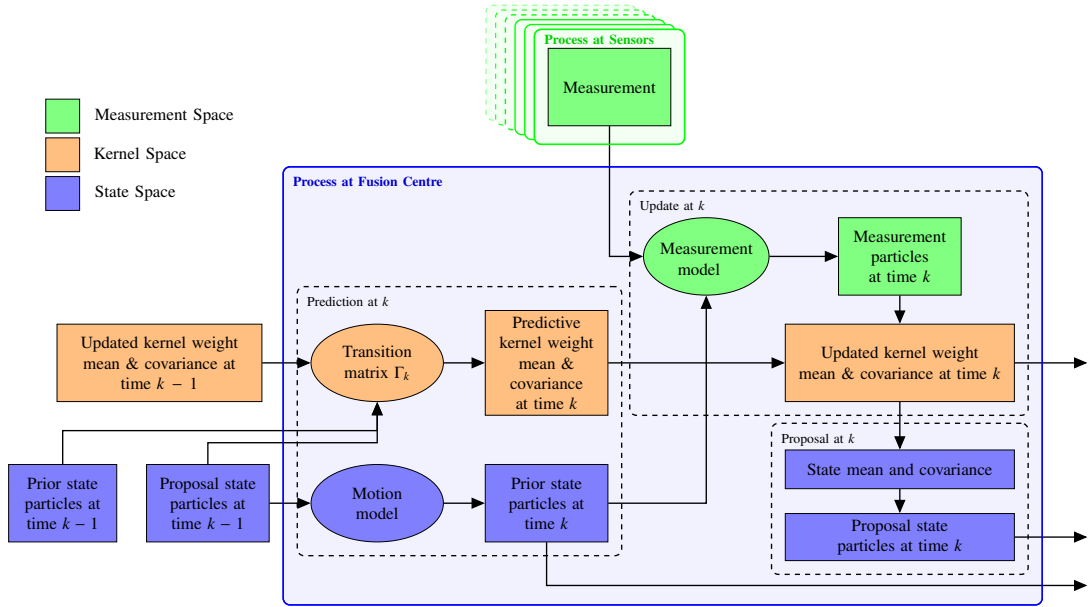
REFERENCES

- [1] Z. Liu, Y. Shang, T. Li, G. Chen, Y. Wang, Q. Hu, and P. Zhu, "Robust multi-drone multi-target tracking to resolve target occlusion: A benchmark," *IEEE Transactions on Multimedia*, vol. 25, pp. 1462–1476, 2023.
- [2] J. S. Wright, J. R. Hopgood, M. E. Davies, I. K. Proudler, and M. Sun, "Implementation of adaptive kernel Kalman filter in Stone Soup," in *2023 Sensor Signal Processing for Defence Conference (SSPD)*, 2023, pp. 1–5.
- [3] S. Wang, F. Jiang, B. Zhang, R. Ma, and Q. Hao, "Development of uav-based target tracking and recognition systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 8, pp. 3409–3422, 2020.
- [4] M. Sun, M. E. Davies, I. K. Proudler, and J. R. Hopgood, "Adaptive kernel Kalman filter," *IEEE Transactions on Signal Processing*, vol. 71, pp. 713–726, 2023.
- [5] —, "Adaptive kernel Kalman filter based belief propagation algorithm for maneuvering multi-target tracking," *IEEE Signal Processing Letters*, vol. 29, pp. 1452–1456, 2022.
- [6] S. Hiscocks, O. Harrald, J. Barr, N. Perree, L. Vladimirov, M. Harris, G. Webb, T. Glover, R. Green, O. Rosoman, E. Rogers, I. Dorrington, J. Wright, S. Pike, E. Hunter, B. Fraser, H. Pritchett, J. Osborne, P. Carniglia, C. England, J. Carlson, L. Flaherty, R. Gorman, D. Kirkland, R. Davies, S. Naylor, M. Campbell, D. Last, D. Crew, and C. Sherman, "Stone Soup." [Online]. Available: <https://doi.org/10.5281/zenodo.4663993>
- [7] M. Sun, M. E. Davies, J. R. Hopgood, and I. Proudler, "Adaptive kernel Kalman filter multi-sensor fusion," in *2021 IEEE 24th International Conference on Information Fusion (FUSION)*, 2021, pp. 1–8.
- [8] L. Yang, Z. Liu, X. Wang, and Y. Xu, "An optimized image-based visual servo control for fixed-wing unmanned aerial vehicle target tracking with fixed camera," *IEEE Access*, vol. 7, pp. 68 455–68 468, 2019.
- [9] "Stone Soup's documentation," 2024. [Online]. Available: <https://stonesoup.readthedocs.io/en/latest/>
- [10] S. Hiscocks, J. Barr, N. Perree, J. Wright, H. Pritchett, O. Rosoman, M. Harris, R. Gorman, S. Pike, P. Carniglia, L. Vladimirov, and B. Oakes, "Stone Soup: No Longer Just an Appetiser," in *2023 26th International Conference on Information Fusion (FUSION)*, 2023, pp. 1–8.
- [11] Stone Soup Contributors, "An introduction to Stone Soup: using the Kalman filter," https://stonesoup.readthedocs.io/en/stable/auto_tutorials/01_KalmanFilterTutorial.html.

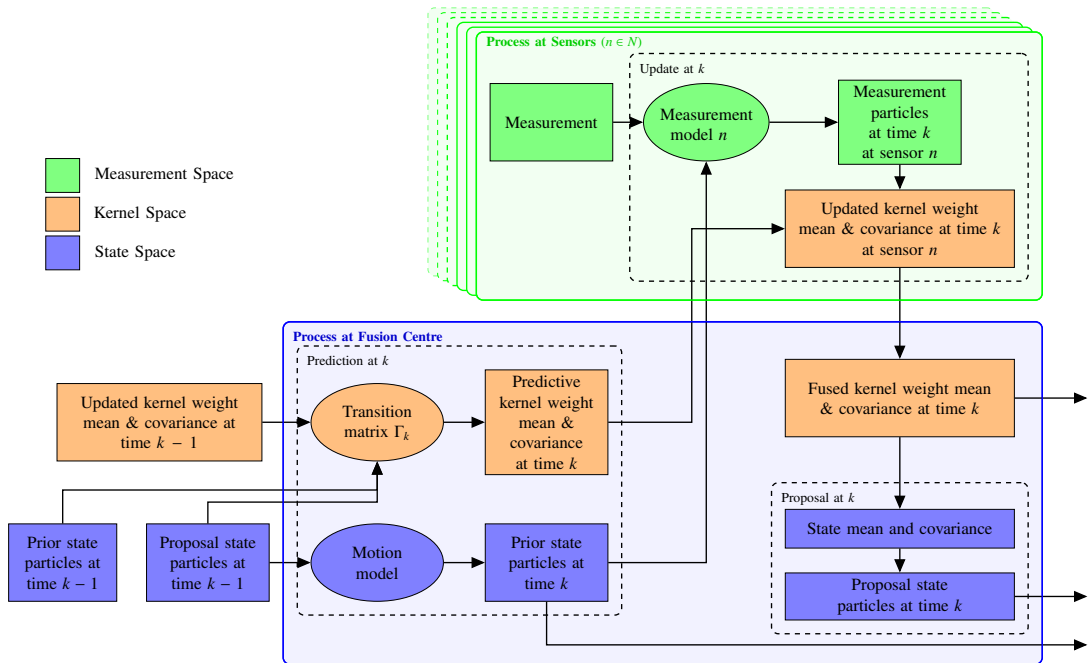
- [12] T. Lemaire, S. Lacroix, and J. Sola, "A practical 3d bearing-only slam algorithm," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005, pp. 2449–2454.
- [13] M. Sun, T. Horton, R. Hodgskin-Brown, M. E. Davies, I. K. Proudler, and J. R. Hopgood, "Adaptive kernel Kalman filter for magnetic anomaly detection-based metallic target tracking," in *2023 Sensor Signal Processing for Defence Conference (SSPD)*, 2023, submitted, pp. 1–5.



(a)



(b)



(c)

Fig. 5: Flow diagrams for (a) the standard AKKF, (b) the centralised Fusion AKKF and (c) the Semi-decentralised Fusion AKKF.