

A Novel Approach to Image Calibration in Super-Resolution Microscopy

Isabel Schlangen, Jérémie Houssineau and Daniel Clark

Abstract—For many disciplines in natural sciences like biology, chemistry or medicine, the invention of optical microscopy in the late 1800's provided groundbreaking insight into biomedical mechanisms that were not observable before with the unaided eye. However, the diffraction limit of the microscope gives a natural constraint on the image resolution since objects which are smaller than half the wavelength of the illuminating light – such as proteins or ions – cannot be recognised in classical microscopy. Recently, different techniques have been developed to partly overcome this restriction using fluorescent molecules as markers. Like this, it is possible to monitor a vast diversity of intracellular processes on a molecular level which are of interest for biomedical research. Since these developments in superresolution microscopy are quite recent, suitable data analysis techniques are still to be advanced. This work aims to deploy the potential of the so-called Hypothesised filter for Independent Stochastic Populations (HISP) for multi-object estimation in a biomedical context by extending its framework to a novel joint object state and sensor drift estimator.

I. INTRODUCTION

During the past decade, super-resolution microscopy has evolved extensively due to its impact on biomedical research. Novel techniques were developed to overcome the diffraction barrier of optical microscopes, either by true sub-wavelength imaging using passive illumination like near-field scanning optical microscopy (NSOM, [1]), Spatially Modulated Illumination (SMI, [2]) and 4Pi microscopy ([3]), or by so-called functional methods, for example PALM (PhotoActivated Localization Microscopy, [4]) and STORM (STochastic Optical Reconstruction Microscopy, [5]). The latter two techniques are based on the idea of active illumination by labeling the molecules of interest with a fluorescent marker and sensing the light that these emit after excitation.

Functional super-resolution imaging techniques in particular have great potential to advance the research of cell-related diseases such as diabetes or cancer since they can be performed on living cells. However, since these methods are quite recent, suitable data analysis techniques have to be developed which can cope with the following imaging artifacts inherent to super-resolution microscopy:

1) *Low SNR*

In order to capture phenomena which last only for nanoseconds, the frame rate has to be as high as possible. However, the amount of photons illuminating the sensor decrease for higher frame rates such that the molecules lose contrast against the background noise.

2) *Bleaching*

After excitation through an external light source, the molecules begin to fluoresce. Over time, the intensity of photon emission decreases which again affects the visibility of the structures of interest.

3) *Drift*

Due to the small size of the monitored objects being only a few nanometers wide, even the smallest movements of the microscope induced by motor vibrations or thermal expansion could have a visible effect on the acquired images. Modern microscopes are equipped with a drift corrector already but they depend on the incorporation of reliable beads which is not always easily provided.

The first two issues can be solved on the image preprocessing level using suitable denoising and image enhancement techniques that will be described below. The correction of sensor drift, however, is usually performed separately from the actual task. Many image registration methods have been developed in the past many of which heavily rely on the extraction of reliable feature points [6]. In some cases, however, the moving objects are the only possible features and stable markers are not always available. For that reason, it is desirable to integrate the estimation of the sensor drift in the multi-object estimation process.

Previous research proved that it is possible to use particle filtering to estimate the non-linear motion of the sensor and formulate the multi-object estimation as a dependent process. In [7], a Simultaneous Localisation and Mapping (SLAM) technique is used to jointly estimate both the trajectory of the robotic vehicle and the motion of its surroundings. Another important task is joint object tracking and sensor calibration [8] which is based on the same idea. Moreover, [9] resp. [10] extend this concept to introduce a first joint molecule and sensor motion estimator for biomedical applications, using the Probability Hypothesis Density (PHD) filter.

The aim of this work is to formulate an alternative to the method of [10] and to show the potential of simultaneous estimation techniques for superresolution microscopy. The new approach will incorporate the Hypothesised filter for Independent Stochastic Populations (HISP) which was formulated by Houssineau, Del Moral and Clark in 2013.

The introduction of the HISP filter gives a whole new perspective on the task of multi-object target tracking. In [11], a novel framework for treating multi-object estimation is introduced which involves the concept of partial distinguishability of objects which has not been considered in previous approaches like Multi-Hypothesis Tracking (MHT) or

Isabel Schlangen has a PhD studentship sponsored by ESRIC (the Edinburgh Super-Resolution Imaging Consortium).

The authors are with the School of Engineering and Physical Sciences (EPS) at Heriot Watt University, Edinburgh EH14 4AS, UK {is117, jh207, D.E.Clark} @hw.ac.uk

PHD filtering. Since first experiments have proved superior performance to the PHD filter especially in cases of low detection rates [12], it is desirable to explore its potential for calibration purposes. Thus, a novel joint multi-object and sensor drift estimator will be introduced on the following, extending the formulation of the general HISP filter.

II. IMAGE PRE-PROCESSING

In [13], different image enhancement techniques for fluorescence microscopy imaging are presented and compared. One of the more lightweight solutions in terms of computational effort is a method based on the so called à trous wavelet decomposition. Similar to the Difference-of-Gaussians method, the input image is convolved with a filter¹ several times and the differences between the outputs are computed.

In particular, let us define the one-dimensional kernel $H = [\frac{1}{16}, \frac{1}{4}, \frac{3}{8}, \frac{1}{4}, \frac{1}{16}]$ and denote the grayscale input image by $\mathcal{I} = \{I(i, j)\}_{i \in \{1 \dots m\}, j \in \{1 \dots n\}} \in \mathbb{R}^{m \times n}$. By sequential row- and column-wise convolution with H , a sequence $\{\mathcal{I}_k\}_{k \in \{1 \dots K\}}$ is obtained where K denotes the number of applied convolution operations. Calculating the differences $\mathcal{W}_k = \mathcal{I}_{k-1} - \mathcal{I}_k$ where $W_k(i, j) = I_{k-1}(i, j) - I_k(i, j)$ for $i \in \{1 \dots m\}$ and $j \in \{1 \dots n\}$, we find the à trous wavelet decomposition

$$\mathcal{I} = \mathcal{I}_K + \sum_{k=1}^K \mathcal{W}_k. \quad (1)$$

Each of the layers \mathcal{W}_k contains features of different sizes, so the image noise can easily be filtered out by thresholding each of the layers before reconstructing the image via (1), replacing \mathcal{W}_k by the thresholded images.

Fig. 1a demonstrates the performance of the algorithm on the example of an input image which was acquired using STORM, showing fluorescently labelled calcium channels in the brain cell membrane of a rat. The resulting image is almost free of noise and shows only the relevant structures in the image. Since the focus of this method is on the size of the structures rather than the absolute intensities, possible bleaching effects are removed automatically.

III. A PHD/HISP CALIBRATION METHOD

As mentioned above, the high resolution of imaging techniques such as PALM and STORM leads to a high sensitivity towards the slightest displacements between the slide and the microscope head, and reliable markers are not easily provided for the drift correction. Thus, it is desirable to estimate the sensor state additionally to the movement of the molecules. In [10], the molecule tracking was performed using PHD filtering which was conditioned on the sensor state. The goal of this paper is to find an extension of the

¹Note that the word *filter* is ambiguous here. In the field of image processing and thus in the current context, filters are masks that introduce a certain effect like smoothing or sharpening on an image when they are convolved. Another class of filters are the Bayesian filters which are used for probabilistic target tracking.

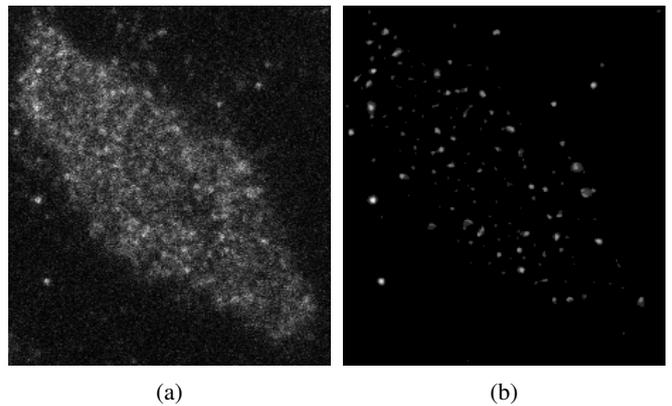


Fig. 1: Output of the à trous wavelet method on an example image acquired using STORM. Fig. 1a shows the noisy input image and Fig. 1b demonstrates the respective performance of the filter.

HISP filter which involves a dependence on the state of the microscope.

Against the usual conventions, we choose to make use of notations from the field of probability theory as proposed in [12]. Thus, denote by $\mathcal{P}(E)$ the set of probability measures on a measurable space $E := (E, \mathcal{E})$ and set $\mu(f) = \int \mu(dx) f(x)$ for any $\mu \in \mathcal{P}(E)$ and an arbitrary bounded measurable function f on E .

In the following, let Ξ_t denote the sensor state space and let ξ_t be the sensor state at a given time t . Moreover, write $p_t^\Xi \in \mathcal{P}(\Xi_t)$ for the sensor state probability measure.

The concept of the HISP filter is to describe targets by sequences of measurements, called observation paths. Therefore, assume that we have a population \mathcal{X} with individuals described in the extended space $\bar{\mathbf{X}}_t = \mathbf{X}_t \cup \{\psi\}$, where \mathbf{X}_t is a complete separable metric space (c.s.m.s.) and ψ denotes the individuals with no image in \mathbf{X}_t . We further define the element φ to describe individuals that are not part of the population \mathcal{X} . Thus, the set of all possible individuals is denoted by $\mathbf{X}_t^+ = \bar{\mathbf{X}}_t \cup \{\varphi\}$.

Observations of individuals in \mathcal{X} are described in the extended observation space $\bar{\mathbf{Z}}_t = \mathbf{Z}_t \cup \{\phi\}$ with \mathbf{Z}_t being a c.s.m.s. and ϕ being the empty observation. Denote the finite set of proper observations z_t at time t by Z_t , and write $\bar{Z}_t = Z_t \cup \{\phi\}$ for the augmented set of observations. Define the Cartesian product $\bar{Y}_t = \times_{i=0}^t \bar{Z}_i$ to be the set of all possible sequences $\mathbf{y}_t := (z_0, \dots, z_t)$ of measurements where each $z_i \in \bar{Z}_i$ was observed at time i . We call the sequences \mathbf{y}_t the observation paths up to time t . Furthermore, define $\phi_t \in \bar{Y}_t$ to be the observation path where nothing has been observed up to time t , i.e. $\phi_t = (\phi, \dots, \phi)$. The objects with observation path ϕ_t are regarded as indistinguishable. Thus, the set of proper observation paths is denoted by $Y_t = \bar{Y}_t \setminus \{\phi_t\}$.

Additionally, a continuous time interval $T \subseteq [0, t]$ is associated with each potential observation path at time t which allows for distinguishing objects that existed at different times. Thus, for an observation path $\mathbf{y} \in Y_{t-1}$ and T as above, let

us define a possible individual of \mathcal{X} by the tuple $\mathbf{x} = (T, \mathbf{y})$, and denote the set of all possible individuals at time t before the observation update by $X_t \subseteq \{T \subseteq [0, t]\} \times \bar{Y}_{t-1}$. Similarly, define $\hat{X}_t \subseteq \{T \subseteq [0, t]\} \times \bar{Y}_t$ to be the set of the possible individuals at time t after the update.

For any measure $p_t^x \in \mathcal{P}(\mathbf{X}_t^+)$, the scalar value $p_t^x(\bar{\mathbf{X}}_t)$ is the prior probability of \mathbf{x} to be an individual of \mathcal{X} at time t , whereas $\hat{p}_t^x(\bar{\mathbf{X}}_t)$ will denote the posterior probability after the measurement update.

Using these notations, we wish to reformulate the findings of [12] by calculating the probabilities conditioned on the sensor state ξ . For this purpose, we use the joint probabilities

$$p_t^x(\cdot, \xi) = p_t^x(\cdot|\xi) p_t^{\Xi}(\xi), \quad (2)$$

$$\text{resp. } \hat{p}_t^x(\cdot, \xi) = \hat{p}_t^x(\cdot|\xi) p_t^{\Xi}(\xi). \quad (3)$$

A. Initialisation

Initially, the sensor state is set to be at the origin since the first time step can be regarded as the reference position without loss of generality. On the other hand, the set X_0 is defined to only contain the individual $\mathbf{x}_0 = (\{0\}, ())$ which means that nothing can be distinguished since no observations have been made previously. The corresponding probability measure $p_0^{\mathbf{x}_0} =: p_0^b \in \mathcal{P}(\mathbf{X}_0^+)$ describes the first appearance of individuals with representation \mathbf{x}_0 , and the cardinality of the set of newborn targets is described by the distribution c_0^b .

B. Time update

The time update of the HISP filter is independent of the sensor movement, thus both can be propagated separately using appropriate Markov kernels $M_{t|t-1}$ from \mathbf{X}_{t-1}^+ to \mathbf{X}_t^+ and $M_{t|t-1}^{\Xi}$ from Ξ_{t-1} to Ξ_t as follows:

$$\begin{aligned} p_t^{\mathbf{x}'}(f) &= \hat{p}_{t-1}^{\mathbf{x}}(M_{t|t-1}(f|\cdot)), \\ p_t^{\Xi}(f) &= \hat{p}_{t-1}^{\Xi}(M_{t|t-1}^{\Xi}(f|\cdot)) \end{aligned}$$

where $\mathbf{x}' = ([t_0, t], \mathbf{y})$ describes the time-propagated version of $\mathbf{x} = ([t_0, t-1], \mathbf{y})$ and f is an arbitrary test function.

Note that depending on the different cases that can occur, the Markov kernel $M_{t|t-1}$ can also be written down more explicitly by

$$\begin{aligned} M_{t|t-1}(dx'|x) &= p_{S,t}(x)m_{t|t-1}(dx'|x), \\ M_{t|t-1}(\psi|x) &= 1 - p_{S,t}(x), \\ M_{t|t-1}(\psi|\psi) &= 1, \quad M_{t|t-1}(\phi|\phi) = 1 \end{aligned}$$

with $x \in \mathbf{X}_{t-1}$, $x' \in \mathbf{X}_t$ and $m_{t|t-1}$ being a Markov kernel from \mathbf{X}_{t-1} to \mathbf{X}_t . The probability $p_{S,t}(x)$ denotes the survival probability for x . Furthermore, the birth of objects can be modelled in the same manner as described above by introducing a new representation $\mathbf{x} = (\{t\}, ())$ with probability measure $p_t^b \in \mathcal{P}(\mathbf{X}_t^+)$ and cardinality distribution c_t^b .

C. Observation update

Concerning the observation update, the sensor state probability measure has to be involved in Bayes' rule in order to respect the dependence of both processes. Let us first introduce some notation to make the data association between the observations and the individuals more transparent.

Let $p_t^{b_z}$ be the probability measure of the newborn individuals at time t , induced by each measurement $z \in Z_t$. The subset of the population \mathcal{X} which contains the newborn objects is denoted by \mathcal{X}_t^b . Likewise, we define the set \mathcal{X}_t^{op} to be the objects without a corresponding real individual in \mathcal{X} inducing false alarms which are associated with the element ψ . The corresponding representations of the elements of \mathcal{X}_t^b and \mathcal{X}_t^{op} are of the form $\mathbf{x}_t^{b_z} = (\{t\}, \mathbf{y}_z)$ and $\mathbf{x}_t^{op_z} = (\{\psi\}, \mathbf{y}_z)$ respectively where \mathbf{y}_z is the concatenation of ϕ_{t-1} with z . We write $X_t^b := \{\mathbf{x}_t^{b_z}\}_{z \in Z_t}$ and $X_t^{op} := \{\mathbf{x}_t^{op_z}\}_{z \in Z_t}$.

The data association is described by the formulation of a suitable mapping between the augmented set $\mathcal{X}_t^+ = \mathcal{X}_t \cup \mathcal{X}_t^b \cup \mathcal{X}_t^{op}$ and the representations in the measurement space as follows.

In one particular data association, the set Z_t can be decomposed into the disjoint sets Z_t^b , Z_t^{op} and Z_t^d which correspond to detected newborn targets, false alarms and detected existing targets respectively. Since only Z_t^d has corresponding elements in X_t , it is necessary to extend the set X_t by two sets X_t^b and X_t^{op} which stand in one-to-one correspondence to Z_t^b resp. Z_t^{op} . Thus, with the set

$$X_t^+ := X_t \cup X_t^b \cup X_t^{op},$$

we can define an injective function $\varsigma : Z_t \hookrightarrow X_t^+$ which describes the data association between the observations and the potential population at time t :

$$\begin{array}{ccccccc} Z_t & = & Z_t^d & \cup & Z_t^b & \cup & Z_t^{op} \\ \downarrow \varsigma & & \downarrow & & \downarrow 1:1 & & \downarrow 1:1 \\ X_t^+ & = & X_t & \cup & X_t^b & \cup & X_t^{op} \end{array}$$

The restriction $\varsigma|_{Z_t^d} : Z_t^d \rightarrow X_t$ can be made bijective by regarding only the subset

$$X_t^d := \{\mathbf{x} \mid \exists z \in Z_t^d : \varsigma(z) = \mathbf{x}\} \subseteq X_t$$

of the representations of the detected individuals, excluding the mis-detected ones.

Thus, an admissible association between the observation set Z_t and the set of potential individuals X_t is defined by the tuple

$$\mathbf{a} = (\varsigma, X_t, X_t^b, X_t^{op}),$$

where $X_t^+ = X_t \cup X_t^b \cup X_t^{op}$ is induced by a particular decomposition of Z_t as described above. The set of all admissible associations \mathbf{a} will be denoted by Adm_t in the following. Note that there is a whole set of maps ς that are compatible with a particular choice of X_t^+ ; it consists of all permutations of possible associations between the detected individuals and the obtained measurements.

Next, the likelihood function g_t can be defined on \mathcal{X}_t , \mathcal{X}_t^b and \mathcal{X}_t^{op} separately: For individuals in \mathcal{X}_t ,

$$\begin{aligned} g_t(z|x, \mathbf{x}) &= p_{D,t}(x) l_t(z|x), \\ g_t(\phi|x, \mathbf{x}) &= 1 - p_{D,t}(x), \\ g_t(\phi|\psi, \mathbf{x}) &= 1, \quad g_t(\phi|\varphi, \mathbf{x}) = 1 \end{aligned}$$

for any $z \in Z_t$, $x \in \mathbf{X}_t$ and $\mathbf{x} \in X_t$ where $l_t(z|x)$ is the likelihood function for the association of z with x and $p_{D,t}(x)$ stands for the probability of detection. Similarly, we write

$$\begin{aligned} g_t(z'|\psi, \mathbf{x}_{fa}) &= \delta_z(z') p_t^{fa,z}, \\ g_t(\phi|\psi, \mathbf{x}_{fa}) &= 1 - p_t^{fa,z} \end{aligned}$$

for the individuals in \mathcal{X}_t^{op} with the probability $p_t^{fa,z}$ for z to be a false alarm and $\mathbf{x}_{fa} = (\emptyset, (\phi_{t-1}, z))$. Finally, the likelihood function for newborn individuals in \mathcal{X}_t^b is defined to be

$$\begin{aligned} g_t(z'|x, \mathbf{x}_b) &= \delta_z(z') l_t(z|x), \\ g_t(\phi|x, \mathbf{x}_b) &= 0 \end{aligned}$$

with $\mathbf{x}_b = ([t], (\phi_{t-1}, z))$. Here, we made the assumption that newborn targets are detected with probability 1 in order to lower the computational cost.

The probability of associating a possible individual $\mathbf{x} \in X_t^+$ with a measurement $z \in Z_t$ is denoted by the probability

$$p_t^{x,z}(f) = p_t^x(f \cdot g_t(z|\cdot, \mathbf{x}));$$

$p_t^{x,z}$ is a probability measure in $\mathcal{P}(X_t^+)$.

The following theorem reformulates Theorem 1 of [12], incorporating the dependence on the sensor state using joint probabilities as in (2) and (3). Its proof is completely analogous to the findings of [12].

Theorem III.1. *With the settings defined above, the posterior $\hat{p}_t^{x,z} \in \mathcal{P}(X_t^+)$ associating any $z \in Z_t$ with any $\mathbf{x} \in X_t^+$ can be written in the following two equivalent ways:*

$$\hat{p}_t^{x,z} = \frac{w(\mathbf{x}, z, \xi) p_t^{\Xi}(\xi) p_t^{x,z}(\cdot|\xi)}{\int \sum_{\mathbf{x}' \in X_t^+} w(\mathbf{x}', z, \xi') p_t^{\Xi}(\xi') p_t^{x',z}(1|\xi') d\xi'} \quad (4)$$

$$= \frac{w(\mathbf{x}, z, \xi) p_t^{\Xi}(\xi) p_t^{x,z}(\cdot|\xi)}{\int \sum_{z' \in Z_t} w(\mathbf{x}, z', \xi') p_t^{\Xi}(\xi') p_t^{x,z'}(1|\xi') d\xi'} \quad (5)$$

where $w(\mathbf{x}, z, \xi)$ is the joint probability for the association of the sets $X_t^+ \setminus \{\mathbf{x}\}$ and $Z_t \setminus \{z\}$.

In the theorem, the term $w(\mathbf{x}, z, \xi)$ evaluates if the association of \mathbf{x} and z is acceptable in the sense that if both would be removed, the remaining individuals could still be matched reasonably with the remaining measurements.

D. Sensor state estimation

In order to evaluate the probability for a particular association \mathbf{a} , consider the distribution

$$P_t^{\mathbf{a}}(B_t^{\mathbf{a}}) = \prod_{\mathbf{x} \in X_t^+} p_t^{x, \varsigma^{-1}(\mathbf{x})}(A_t^{x, \varsigma^{-1}(\mathbf{x})})$$

where $B_t^{\mathbf{a}} = \times_{\mathbf{x} \in X_t^+} A_t^{x, \varsigma^{-1}(\mathbf{x})}$ for a suitable collection of $A_t^{x, \varsigma^{-1}(\mathbf{x})}$ in the Borel set of $\bar{\mathbf{X}}_t$. In the measurement update, the dependence on the sensor state can be included by writing down the joint probability

$$\begin{aligned} \hat{P}_t^{\mathbf{a}}(B_t^{\mathbf{a}}|\xi) \hat{p}_t^{\Xi}(\xi) &= \frac{P_t^{\mathbf{a}}(B_t^{\mathbf{a}}|\xi) p_t^{\Xi}(\xi)}{\sum_{\mathbf{a} \in \text{Adm}_t} P_t^{\mathbf{a}}(1|\xi) p_t^{\Xi}(\xi)} \\ &= \frac{P_t^{\mathbf{a}}(B_t^{\mathbf{a}}|\xi)}{\sum_{\mathbf{a}} P_t^{\mathbf{a}}(1|\xi)} \underbrace{\frac{\sum_{\mathbf{a}} P_t^{\mathbf{a}}(1|\xi)}{\sum_{\mathbf{a}} p_t^{\Xi}(P_t^{\mathbf{a}}(1|\xi))}}_{\Rightarrow(4),(5)} p_t^{\Xi}(\xi) \end{aligned} \quad (6)$$

where the first fraction in the second equality corresponds to the conditional HISP and the second term can be regarded as the likelihood of the sensor state ξ . The clear isolation of the HISP term inside the joint update makes it possible to handle both processes almost separately.

Unfortunately, the estimation of the sensor state cannot be performed easily with a Kalman filter approach since the principle of the HISP filter is to propagate marginal distributions which does not allow for estimating the correlations with the sensor state. Thus, particle filtering seems to be most appropriate to avoid this issue and to regulate the computational effort.

Therefore, let Ξ_t denote the sensor state space at time t and assume a fixed number of N particles of the form $\{(\xi_t^n, w_n)\}_{n \in \{1 \dots N\}}$ where ξ_t^n represents the estimated state and w_n its corresponding weight. Each particle is combined with an HISP filter for the underlying multi-object tracking (a similar approach based on PHD filtering can be found in [10]). In this context, we define $X_{t,n}^+$ to be the possible individuals at time t which correspond to the n th particle. Let us now specify the initialisation, time update and measurement update for the sensor state estimation.

1) *Initialisation:* Since the first input image can be regarded as a reference frame, the sensor states of all particles will be initialised on the origin without loss of generality. Furthermore, no observations have been made in the past, so all HISP filters are initialised as described above.

2) *Time update:* In the temporal update, the particle filter is propagated by sampling from the previous states according to the state transition model of the system, i.e. $\xi_t^n \sim M_{t-1}^{\Xi}(\xi|\xi_{t-1}^n)$ for all $n \in \{1, \dots, N\}$. Since the movement of the sensor and the target movements are independent, the prediction of the HISP filtering can be performed separately.

3) *Update - Resampling:* The update of the particle filter is performed using Importance Sampling with Rejection Correction (see [8]) as described below. After that, the HISP filters are updated in dependence on the resampled sensor states. In particle filtering, a crucial element is to resample the particles according to their respective weights; in the present case, the likelihood function of the individual HISP filters provides a suitable weighting via

$$\hat{w}_n = L(\bar{Z}_t; X_{t,n}^+, \xi_t^n).$$

Here, $L(\bar{Z}_t; X_{t,n}^+, \xi_t^n)$ denotes the likelihood of the set \bar{Z}_t dependent on the hypotheses $X_{t,n}^+$ carried by the n th particle with state ξ_t^n . A possible specific formulation can be found as a consequence of theorem III.1 via

$$\begin{aligned} L(\bar{Z}_t; X_{t,n}^+, \xi_t^n) &= \sum_{x' \in X_{t,n}^+} w(x', z, \xi_t^n) p_t^{\bar{z}}(\xi_t^n) p_t^{x', z}(1|\xi_t^n) \\ &= \sum_{z' \in \bar{Z}_t} w(x, z', \xi_t^n) p_t^{\bar{z}}(\xi_t^n) p_t^{x, z'}(1|\xi_t^n) \end{aligned}$$

which comes directly from the denominator in the HISP calibration update (4) resp. (5).

While the time update is done by sampling from the old particles according to the underlying motion model under the assumption equal weighting, the resampling in the measurement update is performed according to the importance of each particle given by their respective likelihood. In other words, particles with high weights w_n have higher probability to create offspring than the low-weighted estimations. Hence, we can write the particle distribution for the resampling as a weighted sum of Dirac functions

$$\alpha = \sum_{n=1}^N w_n \delta_{\xi_n}.$$

However, it is not recommended to sample from a Dirac function since such a system can diverge quite quickly. Instead, the distribution is made continuous by convolution with a Gaussian kernel which will be called regularisation kernel in the following². In particular, let d be the dimension of the state space Ξ and define K_h to be a Gaussian rescaled by $\frac{1}{h^d}$ with zero mean and the initial particle covariance given by the sensor motion model, d denoting the dimension of the state space. Thus, the resampling distribution becomes

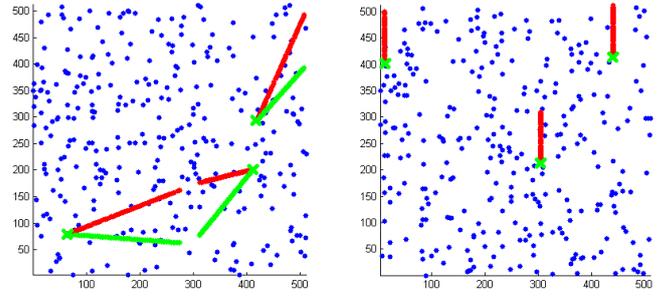
$$\pi(\xi) := \sum_{n=1}^N w_n K_h(\xi - \xi_n).$$

The usage of a regularisation kernel spreads the new particles in the proximity of the respective particle instead of just reproducing it which helps to keep up the diversity of the particle set.

Still, if the weights of the particles are very different from each other, a single resampling from π could still lead to a loss of diversity by only generating new samples from the components with high weights. In [14], a progressive resampling process is proposed which iteratively "flattens" the distribution by potentiation with a factor $c_l < 1$ and resampling therefrom; the series (c_l) slowly converges to 1 over a certain number L of iterations until the original density is regained in the L th step. This method effects in a less aggressive selection of suitable particle positions.

Furthermore, rejection correction is used to prevent being stuck in local maxima, inspired by the approach of Markov Chain Monte Carlo sampling [14]. Thus, a newly sampled

²In fact, other symmetric probability density functions like the Epanechnikov kernel can be used, see [14].



(a) The first scenario involving three targets with constant velocity whose initial value is Gaussian with $\sigma = 1$. (b) The second scenario showing three static targets having small white noise with $\sigma = 0.1$.

Fig. 2: Examples of the two test scenarios for the algorithm evaluation, estimating a linear drift with constant velocity $(0, 1)$ over 100 frames. The ground truth is shown in green, whereas the simulated measurements are plotted in red. Starting points of the tracks are marked with a cross. The blue marks indicate the total background clutter.

particle $\hat{\xi}$ is rejected and set to its previous state ξ with probability

$$\min \left\{ 1, k \cdot \frac{L(\bar{Z}_t; X_t^+, \hat{\xi})}{L(\bar{Z}_t; X_t^+, \xi)} \right\}$$

for a positive constant k which is a tuning parameter of the algorithm.

IV. EXPERIMENTS

In order to evaluate the performance of the proposed approach, it will be contrasted against the PHD based calibration introduced in [10].

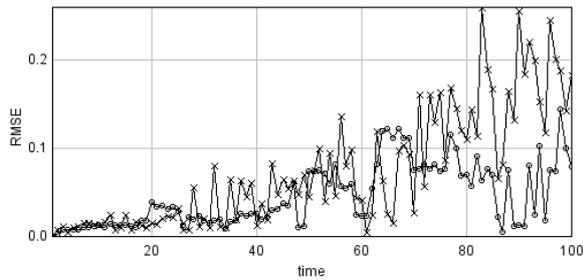
Two scenarios are simulated to highlight different aspects of the calibration. The drift which will be estimated is set to have a constant velocity of $(0, 1)$ throughout all runs, where 1000 particles are used. For the sake of computational complexity, the number of simulated molecules which are present in each frame is set to 3. In both cases, we generate 100 square images of size $512\text{px} \times 512\text{px}$, and independent and identically distributed (i.i.d.) background noise is introduced whose cardinality is Poisson with mean $\lambda = 3$. The probability of detection is set to 0.9.

To cope with the randomness of the particle filter, the results are averaged over 20 Monte Carlo runs.

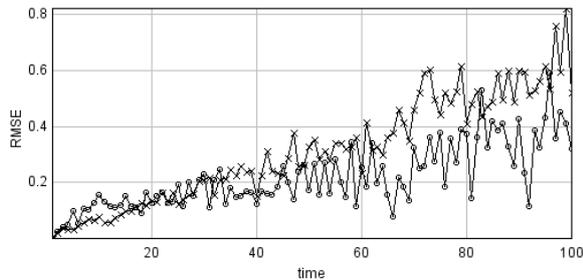
A. First Simulation

The first scenario assumes a linear movement on the three objects in random, but fixed directions up to a small white noise. Their initial positions are set to random, i.i.d. positions in the field of view. An example of the total plot of one Monte Carlo run is shown in Fig. 2a, showing the true positions over all frames in green and the simulated, drifting measurements in red.

In this simulation, the system will be supported by initialising the velocity vector to the true state $(0, 1)$. Like this, it is possible to determine to which extent the movement of the



(a) The drift estimation error in the first scenario over time.



(b) The drift estimation error in the second scenario over time.

Fig. 3: Error plot showing the Root Mean Square Error (RMSE) of the drift position vector over time, averaged over 20 Monte Carlo runs. The crosses mark the PHD filter results, the circles stand for the HISP calibration performance.

targets affects the drift estimation. Resampling is performed with rate $c = 0.8$.

The Root Mean Square Error (RMSE) of the estimated sensor drift over time gives a useful visualisation of the filter’s behaviour. Fig. 3a illustrates that gradually, the uncertainty in the movement of the targets results in growing uncertainty in the drift estimation, yet the error is still reasonably low. However, it can be noted that the HISP filter calibration (marked with \circ) shows an overall better performance, arriving at a mean RMSE value of under 0.1, whereas the PHD calibration error (marked with \times) rises to an average value around 0.2.

B. Second Simulation

In the second case, the three targets are initialised as static objects on random, i.i.d. positions with a small white noise (cf. Fig. 2b). The measurement and motion models for the PHD/HISP filtering are initialised accordingly to the parametrisation of the simulated data, and the resampling rate c is set to 0.2.

In contrast to the first case, the particle positions are initialised at zero position with random velocity, i.e. now, the filter has to estimate the sensor drift without prior knowledge about its true value. Consequently, we set the initial uncertainty of the particle position to 1.1 and the standard deviation of the velocity to 0.1 such that the particle filter has the possibility to explore the state space.

If we plot the RMSE between drift estimate and ground truth against time again (Fig. 3b), it can be seen that the HISP filter outperforms the PHD approach by a factor 2 : 3.

V. CONCLUSION

A novel sensor calibration method has been derived based on the Hypothesised filter for Independent Stochastic Populations (HISP). First experiments showed that it outperforms the existing PHD filter approach already for simple scenarios with high detection rate; even better performance can be expected in scenarios with low detection probability since the classic HISP filter was previously proven to handle low detection rates better than the PHD filter. Thus, the HISP calibration could be of importance for the analysis of super-resolution images in biomedical applications where fiducial markers are not available.

VI. ACKNOWLEDGMENTS

The authors thank the Life Science Interface Laboratory (LSI) at Heriot-Watt for providing the PALM/STORM data.

REFERENCES

- [1] D. W. Pohl, “Optical near field scanning microscope,” 1982, uS Patent No. 4604520.
- [2] M. Hausmann, C. Cremer, J. Bradl, and B. Schneider, “Wave field microscope with detection point spread function,” Mar. 11 2008, uS Patent 7,342,717. [Online]. Available: <http://www.google.com/patents/US7342717>
- [3] S. Hell, “Double-confocal scanning microscope (Doppelkonfokales Rastermikroskop),” 1991, european Patent No. 0491289.
- [4] E. Betzig, G. Patterson, R. Sougrat, O. Lindwasser, S. Olenych, J. Bonifacino, M. Davidson, J. Lippincott-Schwartz, and H. Hess, “Imaging Intracellular Fluorescent Proteins at Nanometer Resolution,” *Science*, vol. 313, No. 5793, pp. 1642–1645, 2006.
- [5] M. Rust, M. Bates, and X. Zhuang, “Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM),” *Nature Methods*, vol. 3, pp. 793–796, 2006.
- [6] B. Zitová and J. Flusser, “Image registration methods: a survey,” *Image and Vision Computing*, vol. 21, pp. 977–1000, 2003.
- [7] C. S. Lee, D. E. Clark, and J. Salvi, “SLAM With Dynamic Targets via Single-Cluster PHD Filtering,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 7,3, pp. 543–552, 2013.
- [8] B. Ristic, D. Clark, and N. Gordon, “Calibration of Multi-Target Tracking Algorithms Using Non-Cooperative Targets,” *13th Conference on Information Fusion (FUSION)*, pp. 1–8, 2010.
- [9] J. Franco Monsalve, “Simultaneous Tracking of Multiple Particles and Sensor Position Estimation in Fluorescence Microscopy Images,” Master Thesis, Heriot-Watt University, Edinburgh, 2013.
- [10] J. Franco, J. Houssineau, D. Clark, and C. Rickman, “Simultaneous Tracking of Multiple Particles and Sensor Position Estimation in Fluorescence Microscopy Images,” *IEEE International Conference on Control, Automation and Information Sciences (ICCAIS)*, pp. 122–127, 2013.
- [11] J. Houssineau, P. Del Moral, and D. E. Clark, “General Multi-Object Filtering and Association Measure,” *5th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 2013.
- [12] J. Houssineau and D. E. Clark, “Hypothesised filter for independent stochastic populations,” 2014, arXiv: 1404.7408.
- [13] I. Smal, M. Loog, W. Niessen, and E. Meijering, “Quantitative Comparison of Spot Detection Methods in Fluorescence Microscopy,” *IEEE Transactions on Medical Imaging*, vol. 29, 2, pp. 282–301, 2010.
- [14] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*, ser. Statistics for Engineering and Information Science. Springer, 2001.