

Dynamic Spike Super-resolution and Applications to Ultrafast Ultrasound Imaging

Giovanni S. Alberti* Habib Ammari† Francisco Romero†
Timothée Wintz‡

March 8, 2018

Abstract

We consider the dynamical super-resolution problem consisting in the recovery of positions and velocities of moving particles from low-frequency static measurements taken over multiple time steps. The standard approach to this issue is a two-step process: first, at each time step some static reconstruction method is applied to locate the positions of the particles with super-resolution and, second, some tracking technique is applied to obtain the velocities. In this paper we propose a fully dynamical method based on a phase-space lifting of the positions and the velocities of the particles, which are simultaneously reconstructed with super-resolution. We provide a rigorous mathematical analysis of the recovery problem, both for the noiseless case and in presence of noise. Several numerical simulations illustrate and validate our method, which shows some advantage over existing techniques.

We then discuss the application of this approach to the dynamical super-resolution problem in ultrafast ultrasound imaging: blood vessels' locations and blood flow velocities are recovered with super-resolution.

Key words. Super-resolution, dynamic spikes, ultrafast ultrasound imaging, fluorescence microscopy, blood flow imaging, total variation regularization.

Mathematics Subject Classification. 65Z05, 42A05, 42A15, 94A08, 94A20, 65J22.

1 Introduction

It is well-known that the resolution of any wave imaging method is limited by the diffraction limit [3]. Super-resolution is understood as any technique

*Department of Mathematics, University of Genoa, Via Dodecaneso 35, 16146 Genova, Italy (alberti@dima.unige.it).

†Department of Mathematics, ETH Zürich, Rämistrasse 101, 8092 Zürich, Switzerland (habib.ammari@math.ethz.ch, francisco.romero@sam.math.ethz.ch).

‡Sony CSL Paris, 6 rue Amyot, 75005 Paris, France (timothee.wintz@sony.com).

whose resolution surpasses this fundamental limit, which is of order of half the operating wavelength. More precisely, the super-resolution problem can be stated as follows: given low frequency measurements of a medium, typically obtained with a convolution by a low pass filter, reconstruct the original medium with a resolution exceeding the diffraction limit. Since high frequencies are completely lost in the measuring process, it is impossible to solve this problem in the general case. It is then natural to focus on the particular case when the medium is made of a finite number of point sources, with unknown locations and intensities. This framework finds applications in many imaging modalities, from the pioneering work on super-resolved fluorescence microscopy [19, 16, 8, 20, 27] (Nobel Prize in Chemistry 2014 [1]) to the works on super-focusing in locally resonant media [5, 6, 4, 21, 22] and the more recent findings in ultrafast ultrasound localization microscopy [15, 18].

In mathematical terms, each point source is represented as a Dirac delta $w_i\delta_{x_i}$, with unknown locations $x_i \in X \subseteq \mathbb{R}^d$ and intensities $w_i \in \mathbb{C}$. We have the sparse spike reconstruction problem: recover

$$\mu = \sum_{i=1}^N w_i \delta_{x_i}$$

from the measurements $y = \mathcal{F}\mu$, where

$$\mathcal{F} : \mathcal{M}(X) \rightarrow \mathbb{R}^n \tag{1}$$

is the measurement operator from the set of Radon measures $\mathcal{M}(X)$ defined on X . Since $\mathcal{M}(X)$ is infinite dimensional, \mathcal{F} is not injective, and therefore one has to use regularization to recover μ . The common choice in this context is an infinite dimensional variant of the ℓ^1 minimization:

$$\min_{\nu \in \mathcal{M}(X)} \|\nu\|_{\text{TV}} \quad \text{subject to } \mathcal{F}\nu = y. \tag{2}$$

Mathematical theory on this problem has greatly advanced over the past years. It includes stable reconstruction of spikes with separation in one and multiple dimensions [11, 10], robust recovery of positive spikes in the case of a Gaussian point spread function with no condition of separation [7], exact reconstruction for positive spikes in a general setting [12] and the corresponding stability properties [14, 25].

Whenever the dynamics of the medium is relevant, as in the case of blood vessel imaging, the super-resolution problem for dynamic point reflectors μ_t arises. In this case, at each time step one measures $\mathcal{F}\mu_t$ and needs to reconstruct both the locations of the spikes and their dynamics: we call this problem the dynamic spike super-resolution problem (see Figure 1). The current approach for this problem is to perform a static reconstruction at each time step (using the method discussed above), and then to track the spikes to obtain their velocities [18]. This approach suffers from three main drawbacks: first, a lot of data are discarded whenever static reconstruction cannot be performed because of

particles being too close, second, the information from neighboring frames is ignored in the first step of the reconstruction and, third, tracking algorithms are computationally expensive.

In this paper we propose a new method for this dynamical super-resolution problem based on a fully dynamical inversion scheme, in which the spikes' locations and velocities are simultaneously reconstructed. After assuming a (local) linear movement of the spikes, we lift the problem to the phase space, in which each spike becomes a particle with location and velocity. The super-resolution problem is then set in this augmented domain, and the minimization performed directly with the full dynamical data. We provide a theoretical investigation of this technique (the analysis shares some common aspects with the one presented in [17] for a similar problem), including exact and stable recovery properties, as well as extensive numerical simulations. These simulations show the great potential of this approach for practical applications, far beyond the predictions of the theory, which we believe can be further developed.


As mentioned above, one of the main motivations and applications of this work is ultrafast ultrasound localization microscopy [15, 18]. Ultrafast ultrasonography is a recent imaging modality based on the use of plane waves instead of the usual focused waves [23, 26, 13]. The resolution of ultrafast ultrasound is determined by the wavelength of the incident wave, and by other factors such as the length of the receptor array and the range of angles used in angle compounding [2]. Due to diffraction theory, the highest resolution is half a wavelength, which is of the order of $300\ \mu\text{m}$. Thus, in blood vessel imaging, blood vessels separated by less than $300\ \mu\text{m}$ cannot be distinguished. As in fluorescence microscopy, randomly activated micro-bubbles in the blood may be used to produce very localized spikes in the observations, giving rise to a dynamical super-resolution problem [18], in which both the locations and the velocities of the bubbles are of interest (the velocities are also used to estimate the thickness of the blood vessels). This is a framework when our approach can be immediately applied.

This paper is structured as follows. In Section 2 we describe the dynamical super-resolution problem and discuss the method introduced in this paper: the phase-space lifting. In Sections 3 and 4 we study the exact recovery issue in absence of noise, while in Section 5 we prove a stability result for noisy measurements. In Section 6 we provide several numerical simulations which validate the method and the theoretical results. In Section 7 we discuss the applications of this technique to ultrafast ultrasound localization microscopy. Finally, Section 8 contains some concluding remarks and future perspectives.


2 Setting the stage

2.1 The space-velocity model

Let us now introduce our model for super-resolution of dynamic spikes. Instead of considering a single measure μ , we consider a time-varying measure μ_t , where



(a) Position and velocity of the particles: the dots correspond to the positions of particles in the middle frame ($t = t_0$), and the arrows to the displacement of the particles from the first frame to the last one.



(b) Corresponding measured sequence $\{y_k\}_k$.

Figure 1: Illustration of the dynamic spike super-resolution problem, with parameters $d = 2$, $N = 2$ and $K = 2$ and \mathcal{F} is a convolution operator by a Gaussian point spread function.

$t \in [-\delta, \delta]$, and $\delta > 0$ defines our observation window. Since δ is expected to be small, we can approximate the dynamics of each point linearly. Therefore we model each point source as a particle displacing with a constant velocity:

$$\mu_t = \sum_{i=1}^N w_i \delta_{x_i + v_i t}, \quad t \in [-\delta, \delta],$$

where $v_i \in \mathbb{R}^d$. The measurement vector is then composed of uniform samples in the observation window $t_k = k\tau$ for $k \in \{-K, -K + 1, \dots, K - 1, K\}$, with $K = \delta/\tau \in \mathbb{N}$ (where \mathbb{N} denotes the set of all positive integers):

$$y_k = \mathcal{F}\mu_{t_k}, \quad k \in [-K, K].$$

Figure 1 illustrates the space-velocity model in two dimensions.

In this work, we show that under certain conditions, we are able to recover the positions x_i , the velocities v_i and the weights w_i simultaneously with infinite precision, using a sparse spike recovery method.

From now on we will assume that the phase space particles, understood as positions x_i and velocities v_i , lie inside the domain

$$\Omega = \{(x, v) \in \mathbb{R}^{2d} : x + k\tau v \in [0, 1]^d, \quad \forall k \in [-K, K]\}, \quad (3)$$

that is the space-velocity domain in which, for all the considered time samples, the locations of the particles stay inside $[0, 1]^d$. Let $T = \{(x_i, v_i)\}_{i=1}^N$ denote the

set of particles. Furthermore, the set of associated weights w_i will be considered in \mathbb{K} , where \mathbb{K} can be either \mathbb{C} or \mathbb{R} .

The measurement operator $\mathcal{F}: \mathcal{M}([0, 1]^d) \rightarrow \mathbb{C}^n$ is assumed to be of the form

$$\mathcal{F}\nu = (\langle \nu, \varphi_l \rangle)_{l=1}^n, \quad \nu \in \mathcal{M}([0, 1]^d),$$

where $\{\varphi_l\}$ is a family of test functions defined on $[0, 1]^d$ and $\langle \nu, \varphi_l \rangle = \int_{[0, 1]^d} \varphi_l d\nu$. Applying \mathcal{F} to the measures $\mu_k := \mu_{t_k}$ for every time step $k \in \{-K, \dots, K\}$ gives

$$\mathcal{F}\mu_k = (\langle \mu_k, \varphi_l \rangle)_{l=1}^n = \left(\sum_{i=1}^N w_i \varphi_l(x_i + k\tau v_i) \right)_{l=1}^n.$$

By construction, these measurements are composed of a vector of size n for each time sample k . We now describe how to express these measurements via an operator defined on measures on the phase space. Consider the measure $\omega \in \mathcal{M}(\Omega)$ and the family of test functions $\varphi_{l,k} \in C(\Omega)$ given by

$$\omega = \sum_{i=1}^N w_i \delta_{x_i, v_i}, \quad \varphi_{l,k}(x, v) = \varphi_l(x + k\tau v),$$

and the measurement operator

$$\mathcal{G}: \mathcal{M}(\Omega) \rightarrow \mathbb{R}^{n \times (2K+1)}, \quad \mathcal{G}\lambda = (\langle \lambda, \varphi_{l,k} \rangle)_{l,k}. \quad (4)$$

With these objects we can write $(\mathcal{F}\mu_k)_k = \mathcal{G}\omega$, and so the measurements are given by

$$y = \mathcal{G}\omega. \quad (5)$$

The dynamical reconstruction problem is now set in the phase space, and consists in the recovery of the sparse measure ω from the measurements (5). Following the approach for sparse spike recovery, we pose this inversion as a total variation (TV) optimization problem, in which we seek to reconstruct positions and velocities simultaneously by minimizing

$$\min_{\lambda \in \mathcal{M}(\Omega)} \|\lambda\|_{\text{TV}} \quad \text{subject to } \mathcal{G}\lambda = y, \quad (6)$$

where the TV norm of a measure $\lambda \in \mathcal{M}(\Omega)$ is defined as

$$\|\lambda\|_{\text{TV}} := \sup \left\{ \int_{\Omega} f d\lambda : f \in C(\Omega), \|f\|_{\infty} \leq 1 \right\}.$$

We will call (6) the *dynamical* recovery, whereas (2) will be called the *static* recovery. The aim of this section is to determine conditions under which exact recovery holds, namely when ω is the unique minimizer of (6). In this way, the recovery of the measure ω from the data y reduces to a convex optimization problem.

2.2 The perfect low-pass case

Instead of studying the general framework outlined so far, in order to highlight the main features of this approach we prefer to focus on the particular case of low-frequency Fourier measurements, which represents a simplified model for many different applications. Thus, the theoretical analysis discussed below refers only to this situation, even though most parts may be extended to the general case of a convolution operator.

The low-frequency Fourier measurements are expressed by the complex sinusoids $\varphi_l(x) = e^{-2\pi i x \cdot l}$ for $l \in \mathbb{Z}^d$ with $\|l\|_\infty \leq f_c$, where $f_c \in \mathbb{N}$ is the highest available frequency for the considered imaging system. The static measurements are given by

$$(\mathcal{F}\nu)_l = \int_{[0,1]^d} e^{-2\pi i x \cdot l} d\nu(x), \quad l \in \{-f_c, \dots, f_c\}^d.$$

With dynamical data, for $k \in \{-K, \dots, K\}$, we have $\varphi_{l,k}(x, v) = e^{2\pi i(x+k\tau v) \cdot l}$ and so

$$(\mathcal{G}\omega)_{l,k} = \int_{\Omega} e^{-2\pi i(x+k\tau v) \cdot l} d\omega(x, v) = \int_{\Omega} e^{-2\pi i(x,v) \cdot (l, k\tau l)} d\omega(x, v). \quad (7)$$

This expression shows that our data consist of low-frequency samples of $2d$ -dimensional Fourier measurements restricted to the d -dimensional subspaces $\{(\xi, k\tau\xi) \in \mathbb{R}^{2d} : \xi \in \mathbb{R}^d\}$ for $k \in \{-K, \dots, K\}$. Thus, our problem is in principle harder than the one considered in [11], in which one measures all low-frequency Fourier coefficients, and not only

$$\{(l, k\tau l) : l \in \mathbb{Z}^d, \|l\|_\infty \leq f_c, k = -K, \dots, K\}, \quad (8)$$

and different from the one considered in [17], in which the Fourier transform is sampled along lines. For a visual representation of this restriction in the case $d = 1$, see Figure 2.

3 Exact recovery in absence of noise

3.1 Dual certificates

As it is standard in convex optimization, it is useful to consider the dual problem to study the exact recovery for (6). In order to do this, we need to introduce the concept of dual certificate. We use the notation $\text{sgn } \mathbb{K}^N = \{\eta \in \mathbb{K}^N : |\eta_i| = 1 \text{ for every } i = 1, \dots, N\}$.

Definition 1 (Dual certificate). Let $T = \{(x_i, v_i)\}_{i=1, \dots, N} \subseteq \Omega$ be a configuration of particles and $\eta \in \text{sgn } \mathbb{K}^N$. A *dual certificate* of the dynamical recovery problem (6) is a function

$$q(x, v) = \sum_{k=-K}^K \sum_{\|l\|_\infty \leq f_c} c_{k,l} e^{i2\pi l \cdot (x+k\tau v)}, \quad (9)$$




Figure 2: The allowed frequencies for the full low-frequency case (left) and in our case, when $d = 1$, $f_c = 3$ and $K = 2$.

where $c_{k,l} \in \mathbb{C}$, obeying

$$\begin{cases} q(x_i, v_i) = \eta_i, & i \in \{1, \dots, N\}, \\ |q(x, v)| < 1, & (x, v) \in \Omega \setminus T. \end{cases} \quad (10)$$

In the following, we shall say that a function of the form (9) has *dynamical form*.

In Figure 3, we present an example of a dual certificate for the dynamical recovery problem, which was computed using a predefined kernel as in [11].

The existence of a dual certificate guarantees exact recovery for the minimization problem (6). More precisely, we have the following result (for a proof, see [11, Proposition A.1]).

Proposition 2. *Suppose that for every $\eta \in \text{sgn } \mathbb{K}^N$ there exists a dual certificate for the dynamical recovery problem, and let $\hat{\omega}$ be a minimizer of (6). Then $\hat{\omega} = \omega$.*

With this proposition in hand, our problem reduces to finding conditions under which dual certificates exist. As mentioned above, the static recovery problem was treated in [11], but their methodology cannot be transferred directly to our case since the static dual certificates are constructed with all low frequency coefficients, whereas in our case we have access only to the frequencies given by the set (8).

The particular structure of functions with dynamical form (9) allows for a simple decomposition

$$q(x, v) = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} q_k(x, v), \quad q_k(x, v) = \sum_{\|l\|_\infty \leq f_c} c_{k,l} e^{i2\pi l \cdot (x + kv)},$$




Figure 3: Example in $d = 1$ of a dual certificate for $N = 5$ particles (inside the red circles) with real weights. The parameters are $K = 1$, $f_c = 20$ and $\tau = 0.025$. The positions, speeds and weights were selected at random.

where $\mathcal{K} \subseteq \{-K, \dots, K\}$ is a subset of the time samples that is used to construct the dual certificate. Observe that the functions q_k are constant along the d -dimensional subspaces parallel to $\{(x, v) \in \mathbb{R}^{2d} : x + k\tau v = 0\}$. Thus, in principle they can be seen as functions of $[0, 1]^d \subseteq \mathbb{R}^d$ instead of $\Omega \subseteq \mathbb{R}^{2d}$. More precisely, we write

$$q_k(x, v) = \tilde{q}_k(x + k\tau v), \quad \tilde{q}_k(y) = \sum_{\|l\|_\infty \leq f_c} c_{k,l} e^{i2\pi l \cdot y}.$$

Consider the values of these functions on the location of the particles at each time

$$\gamma_{i,k} := q_k(x_i, v_i) = \tilde{q}_k(x_i + k\tau v_i), \quad i \in \{1, \dots, N\}, \quad k \in \{-K, \dots, K\}. \quad (11)$$

The functions $q_k(x, v)$ are constant along the affine spaces

$$L_{i,k} := \{(x, v) \in \Omega : (x - x_i) + k\tau(v - v_i) = 0\},$$

which contain the particle i . This implies that the constants $\gamma_{i,k}$ propagate along them, namely

$$q_k(x, v) = \gamma_{i,k}, \quad (x, v) \in L_{i,k}.$$

In other words, the values of the dual certificate on $L_{i,k}$ are completely determined by $\gamma_{i,k}$. Moreover, by (10), these values must satisfy the conditions

$$\frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \gamma_{i,k} = \eta_i, \quad i \in \{1, \dots, N\}. \quad (12)$$

In Figure 4 we present an example in $d = 1$, with three static particles ($N = 3$) and three time measurements ($K = 1$). The values of q are fixed by $\gamma_{i,k}$ on




Figure 4: The geometries of the problem for some simple configurations in one dimension ($d = 1$) with three time measurements ($K = 1$) and $\mathcal{K} = \{-1, 0, 1\}$

each line $L_{i,k}$, and in particular in their points of intersection. As we shall see below, the problematic points are those where several lines (or d -dimensional affine subspaces) intersect, as in the two circled dots in the figure.

A natural way to build dual certificates for the dynamical problem is to consider dual certificates for the static problems at each time step in \mathcal{K} and then to average them. In particular, we make the choice $\gamma_{i,k} = \eta_i$ for every i and k . More precisely, we have the following definition.

Definition 3 (Static average certificate). Take $T = \{(x_i, v_i)\}_{i=1, \dots, N} \subseteq \Omega$. Let $\eta \in \text{sgn } \mathbb{K}^N$ and $\mathcal{K} \subseteq \{-K, \dots, K\}$ with $|\mathcal{K}| \geq 3$. Assume that for every $k \in \mathcal{K}$ there exists a static dual certificate, i.e. there exists $\tilde{q}_k(x) = \sum_{\|l\|_\infty \leq f_c} c_{k,l} e^{i2\pi l \cdot x}$ such that

$$\begin{cases} \tilde{q}_k(x_i + k\tau v_i) = \eta_i, & i \in \{1, \dots, N\}, \\ |\tilde{q}_k(y)| < 1, & y \in [0, 1]^d \setminus \{(x_i + k\tau v_i)\}_{i=1}^N. \end{cases} \quad (13a)$$

We call the function $q(x, v)$ defined as

$$q(x, v) = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} q_k(x, v) = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \tilde{q}_k(x + k\tau v) \quad (13b)$$

a \mathcal{K} -static average certificate.

If static dual certificates exist for every time sample $k \in \mathcal{K}$, we can immediately build a static average certificate $q(x, v)$ by using (13b). By construction, it satisfies

$$\begin{cases} q(x_i, v_i) = \eta_i, & i \in \{1, \dots, N\}, \\ |q(x, v)| \leq 1, & (x, v) \in \Omega \setminus T. \end{cases} \quad (14)$$

This function almost satisfies (10), except that it may happen that $|q(x, v)| = 1$ for some $(x, v) \in \Omega \setminus T$. Take as example the configuration of points given in Figure 4b with $\eta_1 = \eta_2 = \eta_3 = 1$ and $\mathcal{K} = \{-1, 0, 1\}$. The static average certificate will value 1 in each of the particles and by construction $\gamma_{i,k} = 1$. As a consequence, q will have value 1 also in the circled points, in which $|\mathcal{K}| = 3$ lines $L_{i,k}$ intersect, hence breaking condition (10).

In order to ensure that the configuration of particles $\{(x_i, v_i)\}_i$ admits a dual certificate, we characterize these conflictive points.

Definition 4 (Ghost particles). Let $\{(x_i, v_i)\}_{i=1, \dots, N} \subseteq \Omega$ be a configuration of particles and $\mathcal{K} = \{k_1, \dots, k_m\}$ be $m = |\mathcal{K}|$ time samples. A point $(g, w) \in \Omega$ is a *ghost particle* if there exists a set of different indexes $i_1, \dots, i_m \in \{1, \dots, N\}$ such that

$$\bigcap_{p=1}^m L_{i_p, k_p} = \{(g, w)\}.$$

To give the intuition behind the definition of ghost particles, recall that the elements of Ω represent the trajectories of moving objects in $[0, 1]^d$: given a time k , a particle $(x, v) \in \Omega$ describes an object located in $x + kv$. From this we notice that the set $L_{i,k}$ represents all possible moving objects in $[0, 1]^d$ that at time k would be placed at the same location as the particle i , since $x + kv = x_i + kv_i$. Therefore, ghost particles can be understood as possible objects that at every time sample share their location with a given particle. In the example presented in Figure 4b, the highlighted ghost point on top of the particles represents an object moving from left to right, that for $k = -1, 0$, and 1 , would be located in x_1, x_2 and x_3 respectively.

3.2 Main result

We are now ready to state the main result of this section. Several comments on the assumptions are given after the proof.

Theorem 5. *Let $T = \{(x_i, v_i)\}_{i=1, \dots, N}$ be a configuration of N particles, $w \in \mathbb{K}^N$ and $\mathcal{K} \subseteq \{-K, \dots, K\}$ be such that $|\mathcal{K}| \geq 3$. Let*

$$\omega = \sum_{i=1}^N w_i \delta_{(x_i, v_i)} \in \mathcal{M}(\Omega)$$

be the unknown measure to be recovered. Assume that:

- (1) *for every $k \in \mathcal{K}$ and $\eta \in \text{sgn } \mathbb{K}^N$ there exists a static dual certificate $\tilde{q}_k(x)$ satisfying (13a);*
- (2) *and the configuration does not admit ghost particles.*

Then ω is the unique solution of the dynamical recovery problem (6), where \mathcal{G} is given by (7).

Proof. By Proposition 2, it is sufficient to construct a dual certificate for the dynamical recovery problem. Thanks to assumption (1), for every $\eta \in \text{sgn } \mathbb{K}^N$ we can build a \mathcal{K} -static average certificate $q(x, v)$. By (14) and assumption (2) it is enough to prove that

$$|q(x, v)| = 1 \implies (x, v) \in T \cup G,$$

where G is the set of all the ghost particles of the configuration.

Note that for $k_1 \neq k_2$ and any two $i, j \in \{1, \dots, N\}$, the set $L_{i, k_1} \cap L_{j, k_2}$ has at most 1 element. In particular, we notice that $L_{i, k_1} \cap L_{i, k_2} = \{(x_i, v_i)\}$. This observation will be useful below.

Suppose $|q(x, v)| = 1$. Since $q(x, v)$ is defined as an average of terms $q_k(x, v)$, where each of them satisfies $|q_k(x, v)| \leq 1$, a necessary condition for $|q(x, v)| = 1$ is that for every $k \in \mathcal{K}$ we have $|\tilde{q}_k(x + k\tau v)| = |q_k(x, v)| = 1$. By definition that happens exclusively if for every $k \in \mathcal{K}$, $(x, v) \in L_{i, k}$ for some $i \in \{1, \dots, N\}$. In other words, there exists a family of indexes $i_k \in \{1, \dots, N\}$ such that

$$\{(x, v)\} = \bigcap_{k \in \mathcal{K}} L_{i_k, k}.$$

There are two cases: if some of these indexes repeat (i.e. $i_{k_1} = i_{k_2}$, for $k_1 \neq k_2$), then we know that (x, v) must be equal to particle i_{k_1} . In the case none of the indexes i_k repeats, then by definition (x, v) is a ghost particle. \square

3.3 Comments on the hypotheses of Theorem 5

Let us now comment on assumptions (1) and (2), and show why these are easily satisfied. Let us start from assumption (1), namely the existence of static dual certificates.

Remark 6. Take $k \in \{-K, \dots, K\}$. There exists a static dual certificate $\tilde{q}_k(x)$ satisfying (13a) in any of the following situations.

- (a) The particles at time step $k\tau$ are sufficiently separated, namely

$$\|(x_i + k\tau v_i) - (x_j + k\tau v_j)\|_\infty \geq \frac{C_d}{f_c}, \quad i \neq j, \quad (15)$$

and $f_c \geq C'_d$, where $C_d, C'_d > 0$ are constants depending only on the dimension¹ [11].

- (b) The weights w_i are all positive and the particles at time step $k\tau$ are divided into groups, and within each group a minimum separation condition like (15) is satisfied [24] (this holds in the discrete setting).
- (c) The weights w_i are all positive, $d = 1$ and $f_c \geq 2N$ [12]. (It is remarkable that in this case no minimum separation condition is required.)

¹Several bounds are known for these constants, for instance $C_d = 2$ if $d = 1$ and $\mathbb{K} = \mathbb{C}$, $C_d = 1.87$ if $d = 1$ and $\mathbb{K} = \mathbb{R}$, and $C_d = 2.38$ if $d = 2$ and $\mathbb{K} = \mathbb{R}$. More precise estimates may be derived, which yield slightly better constants.

There are reasons to believe that (c) should be sufficient also if $d > 1$ [25], but as far as we are aware a rigorous proof of this fact is still missing.

Let us now turn to assumption (2), and show that it is satisfied almost surely if the particles (x_i, v_i) are chosen uniformly at random.

Proposition 7. *Assume $|\mathcal{K}| \geq 3$. Let $\{(x_i, v_i)\}_{i=1}^N$ be independent random variables, drawn from absolutely continuous distributions μ_i supported in Ω . Then almost surely there are no ghost points.*

Proof. To simplify the notation denote $P_i = (x_i, v_i) \in \Omega$ and $T_N = \{P_i\}_{i=1}^N$. Set $m = |\mathcal{K}|$. Let $G(T_n)$ denote the set of ghost particles of a configuration T_n of n particles. We now define the potential ghost particles $\tilde{G}(T_{n-1})$ of a configuration T_{n-1} of $n-1$ particles. We say that $(g, w) \in \tilde{G}(T_{n-1})$ if there exists a particle $P_n \in \Omega$ such that $(g, w) \in G(T_{n-1} \cup \{P_n\})$, namely

$$\tilde{G}(T_{n-1}) := \bigcup_{P_n \in \Omega} G(T_{n-1} \cup \{P_n\}).$$

By definition of ghost points, we have that if $(g, w) \in \tilde{G}(T_{n-1})$ then there exist $m-1$ distinct time samples $k_1, \dots, k_{m-1} \in \mathcal{K}$ and $m-1$ distinct particles $P_{i_1}, \dots, P_{i_{m-1}} \in T_{n-1}$ such that

$$(g, w) \in \bigcap_{p=1}^{m-1} L_{i_p, k_p}.$$

Notice that $\tilde{G}(T_{n-1})$ is always finite. This stems from the fact that for any $k_1 \neq k_2$ in \mathcal{K} , the intersection of the sets L_{i, k_1}, L_{j, k_2} are singletons for any $i, j \in \{1, \dots, n\}$. Since $m \geq 3$ and T_{n-1} and \mathcal{K} are finite sets, we have a finite number of sets $L_{i, k}$ to intersect, leading to a finite number of elements in $\tilde{G}(T_{n-1})$.

Now we prove that, for any configuration of particles T_{n-1} such that there are no ghost points, the set of new particles P_n that would generate a ghost point has zero measure. More precisely, if $G(T_{n-1}) = \emptyset$ then

$$\mu_n(\{P_n \in \Omega : G(T_{n-1} \cup \{P_n\}) \neq \emptyset\}) = 0. \quad (16)$$

In order to prove this, notice that if $G(T_{n-1}) = \emptyset$ we have

$$G(T_{n-1} \cup \{P_n\}) \subseteq \{(g, w) \in \tilde{G}(T_{n-1}) : \exists k \in \mathcal{K}, (g, w) \in L_{n, k}\}.$$

Thus, since $(g, w) \in L_{n, k}$ means $(g - x_n) + \tau k(w - v_n) = 0$ we obtain

$$\begin{aligned} & \{P_n \in \Omega : G(T_{n-1} \cup \{P_n\}) \neq \emptyset\} \\ & \subseteq \bigcup_{(g, w) \in \tilde{G}(T_{n-1})} \bigcup_{k \in \mathcal{K}} \{P_n \in \Omega : x_n + \tau k v_n = g + \tau k w\}. \end{aligned}$$

Since this is a finite union of affine subspaces of dimension d , it has zero Lebesgue measure. By the absolute continuity of μ_n , we derive (16).

For $n \in \{2, \dots, N\}$, fix a configuration of T_{n-1} particles. Denoting $dP_i = dx_i dv_i$, $f_i = \frac{d\mu_i}{dP_i}$ and $\mathbb{1}_S = 1$ if S is true and $\mathbb{1}_S = 0$ if S is false, we have

$$\begin{aligned} \int_{\Omega} \mathbb{1}_{G(T_n) \neq \emptyset} f_n(P_n) dP_n &= \int_{\Omega} (\mathbb{1}_{G(T_{n-1}) \neq \emptyset} + \mathbb{1}_{G(T_{n-1}) = \emptyset} \mathbb{1}_{G(T_n) \neq \emptyset}) f_n(P_n) dP_n \\ &= \mathbb{1}_{G(T_{n-1}) \neq \emptyset}, \end{aligned}$$

where in the last equality we used (16) and that $\int_{\Omega} f_n(P_n) dP_n = 1$. Using this property $N - m + 1$ times for $n = N, N - 1, \dots, m$ and setting $\mu = \otimes_{i=1}^N \mu_i$, we obtain

$$\begin{aligned} \mu(\{(P_1, \dots, P_N) : G(T_N) \neq \emptyset\}) &= \int_{\Omega} \cdots \int_{\Omega} \mathbb{1}_{G(T_N) \neq \emptyset} \prod_{i=1}^N f_i(P_i) dP_1 \dots dP_N \\ &= \int_{\Omega} \cdots \int_{\Omega} \mathbb{1}_{G(T_{N-1}) \neq \emptyset} \prod_{i=1}^{N-1} f_i(P_i) dP_1 \dots dP_{N-1} \\ &= \dots \\ &= \int_{\Omega} \cdots \int_{\Omega} \mathbb{1}_{G(T_{m-1}) \neq \emptyset} \prod_{i=1}^{m-1} f_i(P_i) dP_1 \dots dP_{m-1} \\ &= 0, \end{aligned}$$

where the last equality follows from $\mathbb{1}_{G(T_{m-1}) \neq \emptyset} = 0$, since there cannot be any ghost particles if there are more time samples than particles. \square

On the other hand, if ghost particles do arise, the conclusion of Theorem 5 may not be true, even if a minimum separation condition is satisfied at all time steps (so that assumption (1) is satisfied). Even though the probability of a random configuration of particles to produce ghost particles is zero, it is worth considering it since the stability of the problem will deteriorate for nearby configurations (see Section 5).

In Figure 5 we provide an example of this case, with three time measurements ($K = 1$), three particles P_1, P_2 and P_3 and three ghost particles G_1, G_2 and G_3 . The configuration is constructed in such a way that, at each time step, the positions of the ghost particles coincide with those of the physical particles, thereby producing the same measurements. In other words, we have $\mathcal{G}(\sum_i \delta_{P_i}) = \mathcal{G}(\sum_i \delta_{G_i})$ and $\|\sum_i \delta_{P_i}\|_{TV} = \|\sum_i \delta_{G_i}\|_{TV}$, and so the minimization problem (6) has multiple solutions.

In the following result, we generalize this observation to more general configurations.

Proposition 8. *Take $w \in \mathbb{R}_+^m$, with $m = |\mathcal{K}| \geq 3$, and let $\{(x_i, v_i)\}_{i=1}^m \subseteq \Omega$ be a configuration of m distinct particles admitting m distinct ghost particles $\{(g_j, w_j)\}_{j=1}^m \subseteq \Omega$. Let $\omega = \sum_{i=1}^m w_i \delta_{(x_i, v_i)}$. Suppose that for every $k \in \mathcal{K}$ and every $i \in \{1, \dots, m\}$ there exists a unique ghost particle (g_j, w_j) in the affine space $L_{i,k}$, i.e.*

$$g_j - x_i + k\tau(w_j - v_i) = 0.$$




Figure 5: In the case of three time measurements, a configuration of points and speeds that allow multiple reconstructions. The lines in the left diagram represent $L_{i,k}$ for each time sample k and particle i . On the right hand side we can observe the relative position of each particle at each time step.

Then the minimization problem in (6) admits infinitely many solutions.

Proof. Consider the measure

$$h = \sum_{i=1}^m \delta_{(x_i, v_i)} - \sum_{j=1}^m \delta_{(g_j, w_j)} \in \mathcal{M}(\Omega).$$

First notice that $h \neq 0$. If we had $h = 0$, each particle would be a ghost particle. By definition of ghost particle, we would have $(g, w) \in L_{i_1, k_1} \cap L_{i_2, k_2}$ for some $k_1 \neq k_2$ in \mathcal{K} , and so

$$(g, w), (x_{i_1}, v_{i_1}) \in L_{i_1, k_1}, \quad (g, w), (x_{i_2}, v_{i_2}) \in L_{i_2, k_2}.$$

Since (x_{i_1}, v_{i_1}) and (x_{i_2}, v_{i_2}) are two different ghost particles, either L_{i_1, k_1} or L_{i_2, k_2} would contain two ghost particles, contradicting the hypotheses.

The measure h is undetectable, since it belongs to the kernel of the operator \mathcal{G} defined in (4), as we now show. We readily compute

$$(\mathcal{G}h)_{l,k} = \langle h, \varphi_{l,k} \rangle = \sum_{i=1}^m \varphi_l(x_i + k\tau v_i) - \sum_{j=1}^m \varphi_l(g_j + k\tau w_j) = 0.$$

Indeed, by our hypothesis on the ghost particles, we have that at every time sample k , for every position $x_i + k\tau v_i \in \Omega$ there exists only one ghost point such that $g_j + k\tau w_j = x_i + k\tau v_i$. Therefore, each term of the first sum cancels out with one term of the second sum, as desired.

For $\beta \in [0, \min_i w_i]$, consider the measure

$$\omega_\beta = \omega - \beta h = \sum_{i=1}^m (w_i - \beta) \delta_{(x_i, v_i)} + \beta \sum_{j=1}^m \delta_{(g_j, w_j)}.$$

Since $\mathcal{G}\omega_\beta = \mathcal{G}\omega$ and $\|\omega_\beta\|_{TV} = \sum_i (w_i - \beta) + \sum_j \beta = \sum_i w_i = \|\omega\|_{TV}$, we obtain that ω_β is a solution to (6) for every $\beta \in [0, \min_i w_i]$. \square

It is worth observing that the non-uniqueness of solutions arises also if there exists a subset of the particles satisfying the conditions of Proposition 8.

4 Other constructions of dynamical dual certificates

In this section we show that the construction of dynamical dual certificates as static average certificates, although natural and efficient, is not the end of the story. In other words, exact recovery may be guaranteed even if assumptions (1) and (2) of Theorem 5 are not satisfied. More precisely, we provide alternative constructions of dynamical dual certificates for configurations that either do not allow static dual certificates (assumption (1)) or have ghost particles (assumption (2)). In particular, the first case shows an advantage of our space-velocity model over applying static reconstructions at each time sample. This aspect will also be investigated in Section 6 below.

4.1 Dual certificates with no static separation condition

The following example of dual certificate is purely numerical, but shows the possibility of constructing a dual certificate in cases in which static reconstructions are expected to fail.

The chosen configuration is presented in Figure 6a, where we consider the one-dimensional case ($d = 1$), five time measurements ($K = 2$), two static particles barely separated enough to allow a reconstruction and a third moving particle. We give positive weights w_i to the static particles, and a negative weight to the moving one. Since this third particle is at each time measurement close to another particle and has a different sign, it is not possible to localize it with a static reconstruction at any point. We recall that f_c is the maximum imaging frequency, and $\frac{1}{2f_c}$ is far below the optimal minimum separation distance (see Remark 6 and [11, Section 5]).

In Figure 6b we can see a dual certificate for this configuration (with $f_c = 20$). To obtain this dual certificate, we followed the construction done in [11] for the two dimensional case, but considering only the frequencies available in our setting given by the set (8).




Figure 6: An example of a dynamical dual certificate without a static minimum separation condition.

4.2 Dual certificates in presence of ghost particles

As we saw in Figure 4b, in the presence of ghost points the static average dual certificate constructed in (13) is not a valid dynamical dual certificate if the values η_i have a constant sign. Indeed, the static average dual certificate will have absolute value equal to 1 in the ghost particles, since the values $\gamma_{i,k}$ on $L_{i,k}$ are simply set to η_i so that $\frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \gamma_{i,k} = \eta_i$ (see Figure 7a). However, by making a slightly different choice for $\gamma_{i,k}$ it is possible to have simultaneously $q(x_i, v_i) = \eta_i$ and $|q(g, w)| < 1$ on every ghost particle, thereby obtaining a valid dual certificate (see Figure 7b).

This construction is formalized in the following result.


Proposition 9. *Take $d = 1$, $f_c \geq 128$, $K = 1$ and $\Delta x \in [\frac{1.87}{f_c}, 1)$. Consider 3 particles with locations $(-\Delta x, 0)$, $(0, 0)$ and $(\Delta x, 0)$. Then for every $\eta \in \text{sgn } \mathbb{K}^3$ there exists a dual certificate.*

Proof. If η is not a constant vector, the static average certificate (whose existence is guaranteed by Remark 6a) given by (13) is a valid dual certificate, and the result is trivial. Without loss of generality, suppose now that $\eta = (1, 1, 1)$. Set $x_{-1} = -\Delta x$, $x_0 = 0$ and $x_1 = \Delta x$.

We construct the dual certificate $q(x, v)$ using the notation of §3.1 with $\mathcal{K} = \{-1, 0, 1\}$. More precisely, we write

$$q(x, v) = \frac{1}{3} \sum_{k=-1}^1 q_k(x, v) = \frac{1}{3} \sum_{k=-1}^1 \tilde{q}_k(x + k\tau v),$$

where the functions \tilde{q}_k are constructed as follows. The quantities $\gamma_{i,k}$ defined in (11) need to satisfy (12), but also to keep the absolute value of q below 1 in the



(a) The static average dual certificate (corresponding to $\varepsilon = 0$) is not a valid dual certificate since it has value 1 in the two ghost particles.

(b) The valid dual certificate constructed as a perturbation of the static average dual certificate by taking $\varepsilon = 0.08$.

Figure 7: Dual certificates for the configuration of particles of Figure 4b.

ghost particles. In view of the geometrical configuration (see Figure 4b), these conditions are:

$$\begin{aligned} \gamma_{i,-1} + \gamma_{i,0} + \gamma_{i,1} &= 3, & i \in \{1, 2, 3\}, \\ |\gamma_{1,1} + \gamma_{2,0} + \gamma_{3,-1}| &< 3, \\ |\gamma_{1,-1} + \gamma_{2,0} + \gamma_{3,1}| &< 3. \end{aligned}$$

To take a family of solutions, for $\varepsilon \in (0, 1)$ set

$$\begin{aligned} \gamma_{i,-1} = \gamma_{i,1} &= 1 - \varepsilon, & i \in \{1, 3\}, \\ \gamma_{i,0} &= 1 + 2\varepsilon, & i \in \{1, 3\}, \\ \gamma_{2,-1} = \gamma_{2,0} = \gamma_{2,1} &= 1. \end{aligned}$$

For each $k \in \mathcal{K}$, let $\tilde{q}_k(x) = \sum_{l=-f_c}^{f_c} c_{k,l} e^{2\pi i l x}$ be the low frequency trigonometric polynomial such that

$$\tilde{q}_k(x_i) = \gamma_{i,k}, \quad \tilde{q}'_k(x_i) = 0, \quad i = -1, 0, 1,$$

constructed in [11], whose existence follows from $\Delta x \geq \frac{1.87}{f_c}$ (the construction works also if $|\gamma_{i,k}| \neq 1$).

The final step is to ensure that $q(x, v) = \sum q_k(x, v)$ satisfies $|q(x, v)| < 1$ for every $(x, v) \in \Omega$ that is not a particle. When $\varepsilon = 0$, \tilde{q}_k is strictly concave in x_i for every i [11]. Furthermore, the map $\varepsilon \mapsto \tilde{q}_k$ is affine, hence every derivative of \tilde{q}_k is continuous in ε . Therefore, since $\tilde{q}'_k(x_i) = 0$, the local concavity of \tilde{q}_k is preserved for ε small and, consequently, the local maxima of \tilde{q}_k are the interpolation points x_i . This is enough to prove that $q(x, v)$ attains local maxima in the particles $(x_i, 0)$ and in the ghost particles $(0, \Delta x)$ and $(0, -\Delta x)$. Finally, by continuity of $q(x, v)$ with respect to ε , there exists an ε sufficiently small such that $|q(x, v)| < 1$ for every element in Ω that is not a particle. \square

The methodology presented in the proof of Proposition 9 can be applied to more general examples and it can be iterated to deal with the ghost particles one by one. The only requirement is to have a direction in which there are no ghost particles, so it can be used as a “sink of mass”. In the example considered, these directions are described by $L_{1,0}$, $L_{3,0}$, $L_{2,-1}$ and $L_{2,1}$.

5 Stable reconstruction with noise

Following the analysis done in [11] for the static case, we present a stability result for the dynamical problem. We will review their setting and adapt it to our case under the same hypotheses: one dimensional case ($d = 1$), discrete setting and a specific type of noise model with bounded total variation norm.

Let us consider a discrete grid $\Omega_{\#}$ with space width $\Delta_x > 0$ and velocity width $\Delta_v > 0$, namely

$$\Omega_{\#} = ((\Delta_x \cdot \mathbb{Z}) \times (\Delta_v \cdot \mathbb{Z})) \cap \Omega,$$

and we will assume throughout this subsection that our particles are located on the grid, i.e. $T \subseteq \Omega_{\#}$ and $\omega \in \mathcal{M}(\Omega_{\#})$. We also recall from [11] the super-resolution factor in space:

$$SRF_x = \frac{1}{\Delta_x f_c},$$

which can be understood as the ratio between the desired space resolution and the permitted resolution given by the diffraction limit.

We consider, instead of (5), the following input noise model:

$$y = \mathcal{G}(\omega + z), \quad \|P_{\mathcal{G}}z\|_{\text{TV}} \leq \delta,$$

where $z \in \mathcal{M}(\Omega_{\#})$ and $P_{\mathcal{G}} = \Delta_x \mathcal{G}^* \mathcal{G}$ stands for the projection over the frequencies described by (8). The reason of this projection is that all the other frequencies are filtered out by the measurement process of \mathcal{G} . We consider a relaxed version of the noiseless problem (6)

$$\min_{\lambda \in \mathcal{M}(\Omega_{\#})} \|\lambda\|_{\text{TV}} \quad \text{subject to} \quad \|\mathcal{G}\lambda - y\|_1 \leq \frac{\delta}{\Delta_x \|\mathcal{G}^*\|_{\ell^1 \rightarrow \mathcal{M}}}, \quad (17)$$

where $\|\mathcal{G}^*\|_{\ell^1 \rightarrow \mathcal{M}}$ is the operator norm of \mathcal{G}^* .

By strengthening the hypotheses of Theorem 5, we obtain the following stability result.

Theorem 10. *Let $f_c \geq 128$, $d = 1$ and $\Delta_x, \Delta_v > 0$ be such that*

$$\Delta_x^2 \leq \frac{K(K+1)}{3} \tau^2 \Delta_v^2. \quad (18)$$

Let $T = \{(x_i, v_i)\}_{i=1, \dots, N} \subseteq \Omega_{\#}$ be a configuration of N particles, $w \in \mathbb{K}^N$ and $\mathcal{K} = \{-K, \dots, K\}$. Let $\omega = \sum_{i=1}^N w_i \delta_{(x_i, v_i)} \in \mathcal{M}(\Omega)$ be the unknown measure to be recovered.

Suppose that

(i) for every $k \in \mathcal{K}$, the minimum separation condition (15) holds;

(ii) and for all $(x, v) \in \Omega \setminus \bigcup_{i=1}^N B_i$ we have

$$\frac{1}{2K+1} \sum_{k=-K}^K \min \left\{ \min_{i \in \{1, \dots, N\}} |x_i - x + k\tau(v_i - v)|^2, \frac{0.1649^2}{f_c^2} \right\} \geq \Delta_x^2, \quad (19)$$

where $B_i = \{(x, v) \in \Omega : |x - x_i| + K\tau|v - v_i| < 0.1649/f_c\}$ are neighborhoods of each particle in T .

Let $\hat{\omega}$ be a minimizer of (17). Then we have the following stability bound for the error:

$$\|\hat{\omega} - \omega\|_1 \leq C(\text{SRF}_x)^2 \delta$$

for some absolute constant $C > 0$.

Remark 11. The stability condition (19) can be understood as an extension of a condition to prevent ghost particles. As we can notice, if (x, v) is a ghost particle, the sum in the left-hand side values 0.

Proof. Take $\eta \in \mathbb{K}^N$ with $|\eta_i| = 1$. Let $q(x, v)$ be the static average dual certificate given in (13), where the each static dual certificate \tilde{q}_k is constructed as in [11] for every $k \in \mathcal{K}$, thanks to assumption (i). With an abuse of notation, set $q = q_{i,j} = q(i\Delta_x, j\Delta_v)$. Let P_T denote the projection onto the space of vectors supported on T , namely, $(P_T q)_{i,j} = q_{i,j}$ if $(i\Delta_x, j\Delta_v) \in T$ and 0 otherwise.

In view of [11, Theorem 1.5], using the same proof we have the following stability bound for the error:

$$\|\hat{\omega} - \omega\|_1 \leq \frac{4\delta}{1 - \|P_{T^c} q\|_\infty}. \quad (20)$$

It remains to bound $\|P_{T^c} q\|_\infty$ in our discrete grid, namely the values of $q_{i,j}$ outside the set of particles T . In order to do so, we shall use the following estimates from [11, Lemma 2.5]:

$$|\tilde{q}_k(t)| \leq 1 - C_1 f_c^2 (t - (x_i + k\tau v_i))^2, \quad \text{when} \quad |t - (x_i + k\tau v_i)| \leq C_2/f_c, \quad (21)$$

and

$$|\tilde{q}_k(t)| \leq 1 - C_1 C_2^2, \quad \text{when} \quad \min_{(x_i, v_i) \in T} |t - (x_i + k\tau v_i)| > C_2/f_c, \quad (22)$$

where $C_1 = 0.3353$ and $C_2 = 0.1649$.

Take $(x_\#, v_\#) \in \Omega_\# \setminus T$, we want to bound the maximum of $|q(x_\#, v_\#)|$. We deal with two cases: when $(x_\#, v_\#)$ is close to some particle, i.e. $(x_\#, v_\#) \in B_i$ for some i , or when it is not.

Fix i such that $(x_\#, v_\#) \in B_i$, then we can write $(x_\#, v_\#) = (x_i + n_x \Delta_x, v_i + n_v \Delta_v)$ with $(n_x, n_v) \in \mathbb{Z} \times \mathbb{Z} \setminus \{(0, 0)\}$. From the definition of B_i , we have that we can use estimate (21) for every $k \in \mathcal{K}$. Therefore

$$\begin{aligned}
|q(x_\#, v_\#)| &= \frac{1}{2K+1} \left| \sum_{k=-K}^K \tilde{q}_k(x_i + k\tau v_i + (n_x \Delta_x + k\tau n_v \Delta_v)) \right| \\
&\leq \frac{1}{2K+1} \sum_{k=-K}^K (1 - C_1 f_c^2 (n_x \Delta_x + k\tau n_v \Delta_v)^2) \\
&= \frac{1}{2K+1} \sum_{k=-K}^K (1 - C_1 f_c^2 (n_x^2 \Delta_x^2 + 2k\tau n_x n_v \Delta_x \Delta_v + k^2 \tau^2 n_v^2 \Delta_v^2)) \\
&= 1 - \frac{C_1 f_c^2}{2K+1} \left(\left((2K+1) n_x^2 \Delta_x^2 + \frac{K(K+1)(2K+1)}{3} \tau^2 n_v^2 \Delta_v^2 \right) \right) \\
&\leq 1 - C_1 f_c^2 \Delta_x^2 (n_x^2 + n_v^2) \\
&\leq 1 - C_1 f_c^2 \Delta_x^2,
\end{aligned}$$

where we used (18) and that $n_x^2 + n_v^2 \geq 1$.

In the case where $(x_\#, v_\#) \notin B_i$ for every i , we take for each $k \in \mathcal{K}$, $i_k = \arg \min_{i \in \{1, \dots, N\}} |x_i - x_\# + k\tau(v_i - v_\#)|$ and $\mathcal{I}_k = x_{i_k} - x_\# + k\tau(v_{i_k} - v_\#)$. Hence, by (22) we have

$$\begin{aligned}
|q(x_\#, v_\#)| &= \frac{1}{2K+1} \left| \sum_{k=-K}^K \tilde{q}_k(x_{i_k} + k\tau v_{i_k} - \mathcal{I}_k) \right| \\
&\leq 1 - \frac{C_1 f_c^2}{2K+1} \sum_{k=-K}^K \left(\mathcal{I}_k^2 \mathbf{1}_{(-\infty, 0]} \left(|\mathcal{I}_k| - \frac{C_2}{f_c} \right) + \frac{C_2^2}{f_c^2} \mathbf{1}_{(0, +\infty)} \left(|\mathcal{I}_k| - \frac{C_2}{f_c} \right) \right) \\
&\leq 1 - \frac{C_1 f_c^2}{2K+1} \sum_{k=-K}^K \min(\mathcal{I}_k^2, C_2^2 / f_c^2) \\
&\leq 1 - C_1 f_c^2 \Delta_x^2,
\end{aligned}$$

where the last inequality is precisely the stability condition (19).

Therefore we have that $\|P_{T^c} q\| \leq 1 - C_1 f_c^2 \Delta_x^2 \leq 1 - C/\text{SRF}_x^2$, for some absolute constant $C > 0$. By inserting this bound into (20), we conclude the proof. \square

6 Numerical simulations

6.1 Methods

Solving minimization problem (6) in all its generality is not an easy task, since it is nonlinear and infinite dimensional. It is possible to use an analogue discrete problem, where the locations and velocities are fixed on a grid whose size

determines the resolution we want to obtain. However, this method becomes intractable for a fine resolution.

In this work, we seek to validate our approach using a reconstruction method in the continuum. In [11], an algorithm with solutions in the continuum is presented in the one dimensional case, but we require a method for higher dimensions. In [9], the authors develop an algorithm to solve the following problem for any linear operator \mathcal{F} from the space of positive punctual measures to \mathbb{R}^n :

$$\min_{\tilde{\mu}} \|\mathcal{F}\tilde{\mu} - Y\|_2^2 \quad \text{subject to } \|\tilde{\mu}\|_{TV} \leq M. \quad (23)$$

Albeit the proposed algorithm is limited to positive weights, this is a realistic expectation in the case of many physical signals, for instance those generated by micro-bubbles in ultrafast ultrasound imaging. Clearly, a minimizer of (23) will be a minimizer of (6) provided that the initial total variation is known.

Lemma 12. *Assume that μ is the unique solution of (6). Then μ is the unique solution of (23) with $M = \|\mu\|_{TV}$.*

Proof. Since (6) admits a unique minimizer, every $\tilde{\mu} \neq \mu$ such that $\mathcal{F}\tilde{\mu} = Y$ verifies $\|\tilde{\mu}\|_{TV} > M$. Therefore, $\tilde{\mu}$ is the unique minimizer of (23). \square

The codes of the simulations of this paper are available at https://github.com/panchoop/dynamic_spike_super_resolution.

6.2 The measurements

We consider the perfect low-pass filter described in Section 2.2 with $d = 1$ as forward measurement operator, where the measured Fourier frequencies are $\{-f_c, \dots, f_c\}$ for some $f_c \in \mathbb{N}$ and the number of time samples are $2K + 1$ with sampling rate τ . The considered parameters for the simulations are:

$$f_c = 20, \quad K = 2, \quad \tau = 0.5. \quad (24)$$

Each simulation contains a random number of particles, between 4 and 10. Furthermore each particle is generated uniformly in Ω and has an associated random weight taken uniformly between 0.9 and 1.1. In total, we made 18,000 simulations.

In order to validate our dynamic spike super-resolution approach, we compare dynamical and static reconstructions and study the stability of our setting. Dynamical reconstruction refers to the recovery of positions and velocities in the space-velocity domain $\Omega \subset \mathbb{R}^2$, whereas by static reconstruction we mean the recovery of the positions of the particles in $[0, 1]$ at some fixed time.

Since we are simulating and reconstructing particles' locations and velocities in the continuum, there are natural associated numerical errors and we require a criterion to call a reconstruction either a success or a failure. Let $\Delta_x, \Delta_v, \Delta_w > 0$ be the maximal accepted errors in the reconstruction of position, velocity and weight, respectively. More precisely, if we consider a single particle $w_i \delta_{(x_i, v_i)}$

and its associated reconstruction $\tilde{w}_i \delta_{(\tilde{x}_i, \tilde{v}_i)}$, then we say that the particle was successfully reconstructed if

$$|x_i - \tilde{x}_i| \leq \Delta_x, \quad |v_i - \tilde{v}_i| \leq \Delta_v, \quad |w_i - \tilde{w}_i| \leq \Delta_w.$$

We consider a configuration successfully reconstructed if, for each particle, there exists a unique successfully reconstructed particle and there are no additional reconstructed particles. These criteria are analogue for static reconstructions.

Notice that in the case of space and velocity, the error values scale with respect to the maximal imaging frequency (f_c in space and $K f_c \tau$ in velocity), thus it is natural to consider the super-resolution factors

$$\text{SRF}_x = \frac{1}{f_c} \frac{1}{\Delta_x} \quad \text{and} \quad \text{SRF}_v = \frac{1}{f_c K \tau} \frac{1}{\Delta_v},$$

which will be set appropriately in each simulation.

Let us introduce the following measure of separation of a configuration of particles $T = \{(x_i, v_i)\}_i \subseteq \Omega$:

$$\Delta_{dyn}(T) = \max_{k \in \{-K, \dots, K\}}^3 \min_{i \neq j} |x_i - x_j + \tau k(v_i - v_j)|, \quad (25)$$

where \max^3 is the third highest element of the set. The quantity Δ_{dyn} represents condition (1) of Theorem 5, for any subset of measurements $\mathcal{K} \subseteq \{-K, \dots, K\}$ with $|\mathcal{K}| = 3$. We will evaluate the simulated reconstructions against this measure of separation, which will be scaled by $\frac{1}{f_c}$, as the theoretical allowed resolution depends on this value (see Remark 6a). Condition (2) of Theorem 5 on the absence of ghost particles could also be considered, but a formula for that purpose is extremely complicated and in view of Proposition 7 we believe it is not necessary to include it for the following analysis.

6.3 Results

6.3.1 Comparison of dynamical and static reconstructions

For a given configuration of particles, we consider three reconstruction procedures.

- The dynamical reconstruction: we take the whole data and recover both positions and velocities of all particles.
- The static reconstruction: we perform static reconstruction of the positions at each time step independently, and we call it a success if the reconstruction is successful for at least one time step.
- The static 3 reconstruction: we perform static reconstruction of the positions at each time step independently, and we call it a success if the reconstruction is successful for at least three time steps. The rationale of this case is that with three successful static reconstructions, it would be possible to use some tracking technique to reconstruct the velocities afterwards. This case also encodes the theoretical limit given by Theorem 5.




Figure 8: Successful reconstruction rates for the cases described in Section 6.3.1, with $\text{SRF}_x = \text{SRF}_v = 1000$ and $\Delta_w = 0.01$. In the horizontal axis we consider the measure of separation (25) scaled by $1/f_c$.

In Figure 8 we present the rates of successful reconstructions for each of the three cases for randomly generated particles. We observe that the dynamical reconstruction has a much higher reconstruction rate than the static reconstructions for small values of Δ_{dyn} , namely for configurations of particles that are never well-separated. This gives a big advantage of dynamical over static reconstructions, and shows that the assumptions of Theorem 5 are in fact too strict: in practice, successful recovery happens much more often than the current theory predicts. For larger values of Δ_{dyn} , the dynamical reconstruction rate remains slightly below 1: this could be explained by the presence of ghost particles, or by numerical issues of the minimization algorithm.

6.3.2 Robustness to noise

We now study the stability of the dynamical reconstruction method and compare it to the stability for the static approaches. We consider a measurement noise model, given by a normally distributed noise scaled by a factor $\alpha \geq 0$. More precisely, for a frequency l and time sample k , our measurements are

$$(\mathcal{G}\omega)_{l,k} = \sum_{i=1}^N w_i e^{-i2\pi l(x_i + k\tau v_i)} + \alpha (\mathcal{N}_{l,k,1} + i\mathcal{N}_{l,k,2}),$$

where $\mathcal{N}_{l,k,j}$ are independent standardized normal random variables. In Figure 9a we plot the reconstruction rate for different values of α . To understand the noise level, we remind that $w_i \in (0.9, 1.1)$. The desired super-resolution factors and the weight threshold are larger in these experiments, since in the presence of noise we do not expect infinitely resolved reconstructions.

We observe that the dynamical reconstruction method is stable to measurement noise, and this stability is similar to the one of the static reconstruction, as we can see in Figure 9b.




Figure 9: Reconstruction rates for the measurement noise model described in Section 6.3.2, with $\text{SRF}_x = \text{SRF}_v = 40$ and $\Delta_w = 0.05$.

6.3.3 Robustness to curvature of trajectory

We study how the reconstruction algorithm fares when instead of a constant velocity, we consider a curved trajectory for the imaged particles. For this purpose we consider the following dynamics for a particle δ_{x_i, v_i} :


$$x(t) = x_i + v_i t + \frac{a}{2} t^2 = x_i + v_i t \left(1 + \frac{a}{2v_i} t\right).$$

We consider $\beta = \frac{a}{2v_i} \tau K$ as measure of curvature. In Figure 10a we present one example of a considered curvature for a trajectory, and in Figure 10b we present the dynamical reconstruction rates for different values of β .

For the presented examples, we relaxed the super resolution factors to 1, since the algorithm is reconstructing some position and velocity, but reasonably it is not located exactly in the target (x_i, v_i) ; similarly, we relaxed the weight threshold condition. Nonetheless, we notice that the stability of the method with respect to changes in the curvature of the trajectories is quite poor. This may be explained by a not optimal choice of the parameters of the minimization algorithm (which should take into account that the forward model is not exact) and by the difficulties for a linear model to capture nonlinear movements. In any case, higher order models are expected to solve this issue and represent an interesting direction for future research.

7 Applications to ultrafast ultrasonography

In this section, we describe a protocol to apply our method to the problem of super-resoluted imaging of blood vessels arising from ultrafast ultrasonography, as mentioned in the Introduction. The setting is the following: we have a sequence of images of a medium containing blood vessels, in which point reflectors



(a) Comparison of the trajectories of a particle with speed 0.25 and curvature $\beta = 0$ and $\beta = 0.03$.

(b) Successful dynamical reconstruction rates for different levels of curvature.

Figure 10: Reconstruction rates with the curvature model described in Section 6.3.3, with $\text{SRF}_x = \text{SRF}_v = 1$ and $\Delta_w = 0.2$.

are randomly activated during a small time interval. These reflectors are moving inside blood vessels and their velocity are approximately the same as that of blood in the blood vessels.

We assume that we can filter out clutter signal coming from other sources than these reflectors. The recorded images are then convolutions of these point sources by the point spread function (PSF). The PSF of ultrafast ultrasound imaging was derived in [2], but filtering out clutter signal changes the shape of the PSF, which is here approximated by a Gaussian function. However, the analysis of [2] allows for a precise derivation, which we leave for future investigation.

7.1 Fully automated imaging protocol

The blood vessels are not necessarily straight lines, and we therefore restrict to a few frames during which the movement of the point reflectors can be approximated by a straight line. We must also make sure that the reconstruction algorithm works in this sequence, and that reflectors do not appear or disappear in a chosen sequence. Given all these remarks, we propose the following imaging procedure, which fully automatically produces a super-resoluted image of blood vessels with velocities. It is composed of two steps:

Choosing reconstruction intervals. In order to apply our dynamical reconstruction algorithm to this problem, we have to select consecutive frames during which the particles do not appear or vanish from one frame to another. One way to ensure this condition is to select time intervals during which the ℓ^2 norm of the observations is constant. Particles appearing or vanishing make the ℓ^2 norm jump, whereas close particles will make the ℓ^2 norm vary slowly.




Figure 11: Graph of the ℓ^2 norm of the simulated measurements in each frame.

Figure 11 illustrates the variation of ℓ^2 norm as a function of frame number (dashed line), and the selected intervals (red solid line), for the case presented in the numerical experiments below.

Reconstructing position and velocities. We then propose to reconstruct positions and velocities using the algorithm presented in the previous sections, using the PSF of ultrafast ultrasound, in each of the intervals chosen in the previous step. By aggregating all positions obtained using this algorithm, we obtain a super-resoluted image of the blood vessels.

7.2 Numerical experiments

7.2.1 Setting

In order to test this procedure, we generate images using a toy example, in which we have two vertical blood vessels that are slightly separated, with particles flowing upwards and downwards in their respective vessel. The speed of the particles depends linearly on their position and their weights are set to 1. The considered parameters are such that they are similar to the experimental ones:

- Domain size: 1 mm \times 1 mm.
- Pixel size: 0.04 mm \times 0.04 mm.
- Vessel separation: 0.04 mm.
- Point spread function: $(x, y) \mapsto e^{-(x^2+y^2)/2\sigma^2}$, $\sigma = 0.04$.
- Maximum and minimum particles' speeds: 15 mm/s, 5 mm/s.
- Sampling rate: $\tau = 0.002 \text{ s}^{-1}$
- Total acquisition time: 2 s.





Figure 12: Three different measurements with ultrafast ultrasound of the simulated particles.

Since in ultrafast ultrasound microbubble imaging the particles are activated and deactivated randomly, we simulate this behavior in the following fashion. The activation of a single particle is modeled as a Bernoulli random variable on each time sample, whereas the deactivation time is modeled as a Poisson random variable. Further, we include measurement noise as in Section 6.3.2, with $\alpha = 0.01$.

7.2.2 Results

In Figure 11 we present the ℓ^2 norms of the measurements at each time step. As an illustration, in Figure 12 we show three frames of the simulated measurements. In Figure 13a we present the B-mode image, i.e. the average signal intensity over all the frames, which shows that the resolution does not allow for separation of the two vessels. Figure 13b present the reconstructed position and velocities using the algorithm described in this work. The results show that the procedure can reconstruct the positions and the velocities very accurately.



(a) B-mode image of the vessels, taken over 2 seconds of measurements.

(b) Super resolved reconstruction of the simulated particles; the colors represent the velocity in the vertical direction.

Figure 13: Simulated vessels' reconstructions.

8 Conclusion

In this paper, we have introduced and studied a new framework for dynamical super-resolution imaging, which allows for super-resolved recovery of positions and velocities of particles from low-frequency measurements. The presented theoretical results are validated by extensive simulations, related to low-frequency one-dimensional Fourier measurements and to two-dimensional ultrafast ultrasound localization microscopy.

In fact, the numerical experiments show that this approach works much better than what the current theory predicts, and so there is a need of further theoretical investigation. Indeed, the arguments presented in this paper are still based on the static reconstruction, while ideally one should consider the dynamical problem directly. Further, it would be nice to generalize this method to higher order models, in order to relax the assumption of linear trajectories.

References

- [1] The Nobel Prize in Chemistry 2014. http://www.nobelprize.org/nobel_prizes/chemistry/laureates/2014/. Nobel Media AB 2014. Web. 14 Feb 2018.
- [2] Giovanni S. Alberti, Habib Ammari, Francisco Romero, and Timothée Wintz. Mathematical analysis of ultrafast ultrasound imaging. *SIAM J. Appl. Math.*, 77(1):1–25, 2017.
- [3] Habib Ammari. *An introduction to mathematics of emerging biomedical imaging*, volume 62. Springer, 2008.
- [4] Habib Ammari, Brian Fitzpatrick, David Gontier, Hyundae Lee, and Hai Zhang. Sub-wavelength focusing of acoustic waves in bubbly media. *Proc. R. Soc. A*, 473(2208):20170469, 2017.
- [5] Habib Ammari and Hai Zhang. A mathematical theory of super-resolution by using a system of sub-wavelength helmholtz resonators. *Communications in Mathematical Physics*, 337(1):379–428, 2015.
- [6] Habib Ammari and Hai Zhang. Super-resolution in high-contrast media. *Proc. R. Soc. A*, 471(2178):20140946, 2015.
- [7] T. Bendory. Robust recovery of positive stream of pulses. *IEEE Transactions on Signal Processing*, 65(8):2114–2122, April 2017.
- [8] Eric Betzig, George H. Patterson, Rachid Sougrat, O. Wolf Lindwasser, Scott Olenych, Juan S. Bonifacino, Michael W. Davidson, Jennifer Lippincott-Schwartz, and Harald F. Hess. Imaging intracellular fluorescent proteins at nanometer resolution. *Science*, 2006.

- [9] Nicholas Boyd, Geoffrey Schiebinger, and Benjamin Recht. The alternating descent conditional gradient method for sparse inverse problems. In *Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2015 IEEE 6th International Workshop on*, pages 57–60. IEEE, 2015.
- [10] Emmanuel J. Candès and Carlos Fernandez-Granda. Super-resolution from noisy data. *J. Fourier Anal. Appl.*, 19(6):1229–1254, 2013.
- [11] Emmanuel J Candès and Carlos Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on Pure and Applied Mathematics*, 67(6):906–956, 2014.
- [12] Yohann De Castro and Fabrice Gamboa. Exact reconstruction using beurling minimal extrapolation. *Journal of Mathematical Analysis and applications*, 395(1):336–354, 2012.
- [13] Charlie Demene, Thomas Deffieux, Mathieu Pernot, Bruno-Felix Osman-ski, Valerie Biran, Jean-Luc Gennisson, Lim-Anna Sieu, Antoine Bergel, Stephanie Franqui, Jean-Michel Correas, et al. Spatiotemporal clutter filtering of ultrafast ultrasound data highly increases doppler and fultrasound sensitivity. *Medical Imaging, IEEE Transactions on*, 34(11):2271–2285, 2015.
- [14] Quentin Denoyelle, Vincent Duval, and Gabriel Peyré. Support recovery for sparse super-resolution of positive measures. *J. Fourier Anal. Appl.*, 23(5):1153–1194, 2017.
- [15] Yann Desailly, Olivier Couture, Mathias Fink, and Mickael Tanter. Sono-activated ultrasound localization microscopy. *Applied Physics Letters*, 103(17):174107, 2013.
- [16] Robert M. Dickson, Andrew B. Cubitt, Roger Y. Tsien, and W. E. Moerner. On/off blinking and switching behaviour of single molecules of green fluorescent protein. *Nature*, 388:355–358, 1997.
- [17] Charles Dossal, Vincent Duval, and Clarice Poon. Sampling the Fourier transform along radial lines. *SIAM J. Numer. Anal.*, 55(6):2540–2564, 2017.
- [18] Claudia Errico, Juliette Pierre, Sophie Pezet, Yann Desailly, Zsolt Lenkei, Olivier Couture, and Mickael Tanter. Ultrafast ultrasound localization microscopy for deep super-resolution vascular imaging. *Nature*, 527(7579):499–502, 2015.
- [19] Stefan W. Hell and Jan Wichmann. Breaking the diffraction resolution limit by stimulated emission: stimulated-emission-depletion fluorescence microscopy. *Opt. Lett.*, 19(11):780–782, Jun 1994.

- [20] Samuel T Hess, Thanu PK Girirajan, and Michael D Mason. Ultra-high resolution imaging by fluorescence photoactivation localization microscopy. *Biophysical journal*, 91(11):4258–4272, 2006.
- [21] Fabrice Lemoult, Nadège Kaina, Mathias Fink, and Geoffroy Lerosey. Wave propagation control at the deep subwavelength scale in metamaterials. *Nature Physics*, 9(1):55, 2013.
- [22] Geoffroy Lerosey, Julien De Rosny, Arnaud Tourin, and Mathias Fink. Focusing beyond the diffraction limit with far-field time reversal. *Science*, 315(5815):1120–1122, 2007.
- [23] G. Montaldo, M. Tanter, J. Bercoff, N. Benech, and M. Fink. Coherent plane-wave compounding for very high frame rate ultrasonography and transient elastography. *Ultrasonics, Ferroelectrics, and Frequency Control, IEEE Transactions on*, 56(3):489–506, March 2009.
- [24] Veniamin I. Morgenshtern and Emmanuel J. Candès. Super-resolution of positive sources: the discrete setup. *SIAM J. Imaging Sci.*, 9(1):412–444, 2016.
- [25] Clarice Poon and Gabriel Peyré. Multi-dimensional sparse super-resolution. *arXiv preprint arXiv:1709.03157*, 2017.
- [26] M. Tanter and M. Fink. Ultrafast imaging in biomedical ultrasound. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 61(1):102–119, January 2014.
- [27] Michael A Thompson, Matthew D Lew, and WE Moerner. Extending microscopic resolution with single-molecule imaging and active control. *Annual review of biophysics*, 41:321–342, 2012.