Efficient coherent diffractive imaging for sparsely varying objects

Yoav Shechtman,[1,*] Yonina C. Eldar,[2] Oren Cohen,[1] and Mordechai Segev[1]

[1]*Physics Department, Technion - Israel Institute of Technology, Haifa 32000, Israel*
[2]*Electrical Engineering Department, Technion - Israel Institute of Technology, Haifa 32000, Israel*
[*]*joe@tx.technion.ac.il*

Abstract: We demonstrate an efficient phase-retrieval method, allowing the recovery of features of dynamic objects that are sparsely varying – i.e. when the difference between two consecutive frames is small. The method uses redundancy in information between similar consecutive frames to recover the features more robustly than current phase-retrieval methods, and necessitates a considerably smaller number of measurements. Both of these features directly lead to shorter acquisition time, paving the way to coherent diffractive imaging of fast dynamic processes. Numerical simulations show a possible 100-fold improvement in temporal resolution over existing methods.

© 2013 Optical Society of America

OCIS codes: (110.1650) Coherence imaging; (100.3010) Image reconstruction techniques.

References and links

1. M. M. Seibert, T. Ekeberg, F. R. N. C. Maia, M. Svenda, J. Andreasson, O. Jönsson, D. Odić, B. Iwan, A. Rocker, D. Westphal, M. Hantke, D. P. DePonte, A. Barty, J. Schulz, L. Gumprecht, N. Coppola, A. Aquila, M. Liang, T. A. White, A. Martin, C. Caleman, S. Stern, C. Abergel, V. Seltzer, J. M. Claverie, C. Bostedt, J. D. Bozek, S. Boutet, A. A. Miahnahri, M. Messerschmidt, J. Krzywinski, G. Williams, K. O. Hodgson, M. J. Bogan, C. Y. Hampton, R. G. Sierra, D. Starodub, I. Andersson, S. Bajt, M. Barthelmess, J. C. Spence, P. Fromme, U. Weierstall, R. Kirian, M. Hunter, R. B. Doak, S. Marchesini, S. P. Hau-Riege, M. Frank, R. L. Shoeman, L. Lomb, S. W. Epp, R. Hartmann, D. Rolles, A. Rudenko, C. Schmidt, L. Foucar, N. Kimmel, P. Holl, B. Rudek, B. Erk, A. Hömke, C. Reich, D. Pietschner, G. Weidenspointner, L. Strüder, G. Hauser, H. Gorke, J. Ullrich, I. Schlichting, S. Herrmann, G. Schaller, F. Schopper, H. Soltau, K. U. Kühnel, R. Andritschke, C. D. Schröter, F. Krasniqi, M. Bott, S. Schorb, D. Rupp, M. Adolph, T. Gorkhover, H. Hirsemann, G. Potdevin, H. Graafsma, B. Nilsson, H. N. Chapman, and J. Hajdu, "Single mimivirus particles intercepted and imaged with an X-ray laser," Nature **470**(7332), 78–81 (2011).
2. J. J. Turner, X. Huang, O. Krupin, K. A. Seu, D. Parks, S. Kevan, E. Lima, K. Kisslinger, I. McNulty, R. Gambino, S. Mangin, S. Roy, and P. Fischer, "X-ray diffraction microscopy of magnetic structures," Phys. Rev. Lett. **107**(3), 033904 (2011).
3. Y. Suzuki, A. Takeuchi, H. Takano, and H. Takenaka, "Performance test of fresnel zone plate with 50 nm outermost zone width in hard X-ray region," Jpn. J. Appl. Phys. **44**(4A), 1994–1998 (2005).
4. H. Mimura, S. Handa, T. Kimura, H. Yumoto, D. Yamakawa, H. Yokoyama, S. Matsuyama, K. Inagaki, K. Yamamura, Y. Sano, K. Tamasaku, Y. Nishino, M. Yabashi, T. Ishikawa, and K. Yamauchi, "Breaking the 10 nm barrier in hard-X-ray focusing," Nat. Phys. **6**(2), 122–125 (2010).
5. D. Sayre, *Imaging Processes and Coherence in Physics*, Springer Lecture Notes in Physics, vol. 112 (Springer-Verlag, 1980), pp. 229–235.
6. J. Miao, P. Charalambous, J. Kirz, and D. Sayre, "Extending the methodology of X-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens," Nature **400**(6742), 342–344 (1999).
7. J. R. Fienup, "Reconstruction of an object from the modulus of its Fourier transform," Opt. Lett. **3**(1), 27–29 (1978).
8. J. R. Fienup, "Phase retrieval algorithms: a comparison," Appl. Opt. **21**(15), 2758–2769 (1982).
9. A. M. Maallo, P. F. Almoro, and S. G. Hanson, "Quantization analysis of speckle intensity measurements for phase retrieval," Appl. Opt. **49**(27), 5087–5094 (2010).
10. J. Miao, R. L. Sandberg, and C. Song, "Coherent X-ray diffraction imaging," IEEE J. Sel. Top. Quantum Electron. **18**(1), 399–410 (2012).
11. H. N. Chapman, A. Barty, M. J. Bogan, S. Boutet, M. Frank, S. P. Hau-Riege, S. Marchesini, B. W. Woods, S. Bajt, W. H. Benner, R. A. London, E. Plönjes, M. Kuhlmann, R. Treusch, S. Düsterer, T. Tschentscher, J. R. Schneider, E. Spiller, T. Möller, C. Bostedt, M. Hoener, D. A. Shapiro, K. O. Hodgson, D. van der Spoel, F. Burmeister, M. Bergh, C. Caleman, G. Huldt, M. M. Seibert, F. R. N. C. Maia, R. W. Lee, A. Szöke, N. Timneanu, and J. Hajdu, "Femtosecond diffractive imaging with a soft-X-ray free-electron laser," Nat. Phys. **2**(12), 839–843 (2006).

12. I. K. Robinson, I. A. Vartanyants, G. J. Williams, M. A. Pfeifer, and J. A. Pitney, "Reconstruction of the shapes of gold nanocrystals using coherent x-ray diffraction," Phys. Rev. Lett. **87**(19), 195505 (2001).
13. S. Eisebitt, J. Lüning, W. F. Schlotter, M. Lörgen, O. Hellwig, W. Eberhardt, and J. Stöhr, "Lensless imaging of magnetic nanostructures by X-ray spectro-holography," Nature **432**(7019), 885–888 (2004).
14. R. L. Sandberg, A. Paul, D. A. Raymondson, S. Hädrich, D. M. Gaudiosi, J. Holtsnider, R. I. Tobey, O. Cohen, M. M. Murnane, H. C. Kapteyn, C. Song, J. Miao, Y. Liu, and F. Salmassi, "Lensless diffractive imaging using tabletop coherent high-harmonic soft-X-ray beams," Phys. Rev. Lett. **99**(9), 098103 (2007).
15. A. Barty, S. Boutet, M. J. Bogan, S. Hau-Riege, S. Marchesini, K. Sokolowski-Tinten, N. Stojanovic, R. Tobey, H. Ehrke, A. Cavalleri, S. Düsterer, M. Frank, S. Bajt, B. W. Woods, M. M. Seibert, J. Hajdu, R. Treusch, and H. N. Chapman, "Ultrafast single-shot diffraction imaging of nanoscale dynamics," Nat. Photonics **2**(7), 415–419 (2008).
16. P. Hockett, C. Z. Bisgaard, O. J. Clarkin, and A. Stolow, "Time-resolved imaging of purely valence-electron dynamics during a chemical reaction," Nat. Phys. **7**(8), 612–615 (2011).
17. L. M. Stadler, C. Gutt, T. Autenrieth, O. Leupold, S. Rehbein, Y. Chushkin, and G. Grübel, "Hard X ray holographic diffraction imaging," Phys. Rev. Lett. **100**(24), 245503 (2008).
18. M. D. Seaberg, D. E. Adams, E. L. Townsend, D. A. Raymondson, W. F. Schlotter, Y. Liu, C. S. Menoni, L. Rong, C.-C. Chen, J. Miao, H. C. Kapteyn, and M. M. Murnane, "Ultrahigh 22 nm resolution coherent diffractive imaging using a desktop 13 nm high harmonic source," Opt. Express **19**(23), 22470–22479 (2011).
19. Y. Kashter, O. Levi, and A. Stern, "Optical compressive change and motion detection," Appl. Opt. **51**(13), 2491–2496 (2012).
20. M. Maletic-Savatic, R. Malinow, and K. Svoboda, "Rapid dendritic morphogenesis in CA1 hippocampal dendrites induced by synaptic activity," Science **283**(5409), 1923–1927 (1999).
21. M. A. Lauterbach, C. K. Ullal, V. Westphal, and S. W. Hell, "Dynamic imaging of colloidal-crystal nanostructures at 200 frames per second," Langmuir **26**(18), 14400–14404 (2010).
22. S. Gazit, A. Szameit, Y. C. Eldar, and M. Segev, "Super-resolution and reconstruction of sparse sub-wavelength images," Opt. Express **17**(26), 23920–23946 (2009).
23. Y. Shechtman, S. Gazit, A. Szameit, Y. C. Eldar, and M. Segev, "Super-resolution and reconstruction of sparse images carried by incoherent light," Opt. Lett. **35**(8), 1148–1150 (2010).
24. H. Wang, S. Han, and M. I. Kolobov, "Quantum limits of super-resolution of optical sparse objects via sparsity constraint," Opt. Express **20**(21), 23235–23252 (2012).
25. A. Beck and Y. C. Eldar, "Sparsity constrained nonlinear optimization: Optimality conditions and algorithms," arXiv:1203.4580 (2012).
26. Y. C. Eldar and S. Mendelson, "Phase retrieval: Stability and recovery guarantees," arXiv:1211.0872 (2012).
27. K. Jaganathan, S. Oymak, and B. Hassibi, "Recovery of sparse 1-D signals from the magnitudes of their Fourier transform," arXiv:1206.1405v1, (2012).
28. H. Ohlsson, A. Y. Yang, R. Dong, and S. S. Sastry, "Compressive phase retrieval from squared output measurements via semidefinite programming," arXiv:1111.6323v3, (2012).
29. S. Mukherjee and C. S. Seelamantula, "An iterative algorithm for phase retrieval with sparsity constaints: Application to frequency-domain optical-coherence tomography," ICASSP 2012.
30. Y. Shechtman, A. Beck, and Y. C. Eldar, "GESPAR: Efficient phase retrieval of sparse signals," arXiv:1301.1018 (2013).
31. B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics* (John Wiley & Sons, 2007).
32. D. Le Gall, "MPEG: A video compression standard for multimedia applications," Commun. ACM **34**(4), 46–58 (1991).
33. D. J. Townsend, P. K. Poon, S. Wehrwein, T. Osman, A. V. Mariano, E. M. Vera, M. D. Stenner, and M. E. Gehm, "Static compressive tracking," Opt. Express **20**(19), 21160–21172 (2012).
34. E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," IEEE Trans. Inf. Theory **52**(2), 489–509 (2006).
35. D. L. Donoho, "Compressed sensing," IEEE Trans. Inf. Theory **52**(4), 1289–1306 (2006).
36. M. Elad, *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing* (Springer Verlag, 2010).
37. M. F. Duarte and Y. C. Eldar, "Structured Compressed Sensing: From Theory to Applications," IEEE Trans. Signal Process. **59**(9), 4053–4085 (2011).
38. Y. Shechtman, Y. C. Eldar, A. Szameit, and M. Segev, "Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing," Opt. Express **19**(16), 14807–14822 (2011).
39. A. Szameit, Y. Shechtman, E. Osherovich, E. Bullkich, P. Sidorenko, H. Dana, S. Steiner, E. B. Kley, S. Gazit, T. Cohen-Hyams, S. Shoham, M. Zibulevsky, I. Yavneh, Y. C. Eldar, O. Cohen, and M. Segev, "Sparsity-based single-shot subwavelength coherent diffractive imaging," Nat. Mater. **11**(5), 455–459 (2012).
40. D. J. Brady, K. Choi, D. L. Marks, R. Horisaki, and S. Lim, "Compressive holography," Opt. Express **17**(15), 13040–13049 (2009).
41. Y. Rivenson, A. Stern, and B. Javidi, "Compressive Fresnel holography," Journal of Display Technology **6**(10), 506–509 (2010).
42. L. Tian, J. Lee, S. B. Oh, and G. Barbastathis, "Experimental compressive phase space tomography," Opt. Express **20**(8), 8296–8308 (2012).
43. D. P. Bertsekas, *Nonlinear Programming (*Athena Scientific, 1999.)
44. J. R. Fienup and C. C. Wackerman, "Phase-retrieval stagnation problems and solutions," JOSA A **3**(11), 1897–1907 (1986).
45. V. Elser, "Phase retrieval by iterated projections," J. Opt. Soc. Am. A **20**(1), 40–55 (2003).

46. M. Guizar-Sicairos, and J. R. Fienup. "Understanding the twin-image problem in phase retrieval," J. Opt. Soc. Am. A. **29**, 2367-2375 (2012).

1. Introduction

The spatial resolution of any standard imaging system is fundamentally bounded by Abbe's diffraction limit, setting the minimal resolvable feature size even in an ideal imaging system to be comparable to one half of the optical wavelength ($\lambda/2$). Therefore, direct optical imaging of an object with nanometric-scale features, e.g. viruses [1], or nanometric magnetic domains in thin magnetic films [2], requires extreme UV or X-ray light. However, focusing a light beam or constructing an imaging system at such wavelengths is very difficult, due to the lack of suitable optical components. Namely, focusing a hard X-ray beam is typically done with a Fresnel zone plate [3] or with focusing mirrors, where in both cases the current focusing ability is fundamentally limited to a spot much larger than the wavelength. For example, at hard X-ray wavelength ($\lambda \sim 0.1$nm) the current resolution limit, in focusing as well as in imaging, is $\sim$10nm, which is $\sim$170 λ [4].

Another avenue to overcome this practical resolution limit for imaging applications, at extreme UV and X-ray wavelengths, is to use coherent diffractive imaging (CDI) [5,6]. X-ray CDI, in its most basic form – plane-wave CDI, consists of illuminating the object with a coherent X-ray plane wave, and measuring the diffracted intensity pattern in the far-field [6]. The diffracted intensity is then algorithmically inverted to obtain the sought object, most commonly by using an iterative phase-retrieval algorithm [7,8]. The loss of information, which is inherent in the measurement process – where only the magnitude of the diffraction pattern is measured and not the phase – is handled algorithmically using prior knowledge about the region where the object resides (the "support"). However, the algorithmic inversion process is challenging, and the prior knowledge it requires for successful reconstruction is often very precise (e.g., on the location of the object, or the support boundaries, etc.). Moreover, all such phase-retrieval algorithms fundamentally require a high signal to noise ratio (SNR), and a large number of measurements in the diffraction plane, a criterion known as the oversampling criterion. Recent techniques suggested improving the phase-retrieval process by use of a phase diffuser to create speckles in the far-field, thereby increasing the spatial variation of the measured far-field intensity [9]. However, the SNR is still the limiting factor for phase-retrieval problems. For a recent review on CDI see [10].

Due to the high SNR and oversampling requirements, the level of signal (light) needed for a high quality recovery is large. This can be obtained by using a single intense pulse of light, e.g. from a free-electron laser [11,1], or through accumulating a large number of relatively weaker X-ray pulses, e.g. from synchrotron radiation [12,13,2], or using light from a high-harmonic generation process [14]. Since an intense X-ray pulse is destructive for certain objects (e.g. proteins), imaging a dynamic process of such an object while it varies in time is challenging: It can either be done by destructive pump-probe time resolved imaging, where identical copies of a dynamic process are imaged with different time delays between the pump and the probe [15,16], or by accumulation of weaker pulses [14,17,18]. However, pulse accumulation naturally prolongs the acquisition time necessary for a single image; currently, the acquisition time is on the scale of tens of seconds, at least [14,17,18] - far from being suitable for real-time dynamic imaging. In the case of lensless imaging (CDI) of magnetic structures [2,13], the exposure time may be shorter ($\sim$0.2 seconds in [2]) – but the exposure time is still limited by SNR. Thus, finding a technique that would work under lower SNR could facilitate faster acquisition, thereby improving the ability to track the formation of magnetic domains with better temporal resolution. For the specific case of imaging the formation of nanometric domains in magnetic films, improving the temporal resolution may allow viewing phase transitions as they occur [13].

Here, we propose and demonstrate an algorithmic method that performs robust phase-retrieval of sparsely varying video, i.e. when the difference between two consecutive frames is small. The proposed method utilizes the similarity between consecutive frames to compensate for the loss of information associated with the measurements. The loss of information is

caused primarily by the lack of phase information, but our method compensates also for low SNR, and allows using a considerably smaller number of measurements than traditional phase-retrieval methods. As such, our method can be used to recover dynamically varying objects using a considerably shorter acquisition time compared to existing algorithms, without requiring multiple identical samples. Using the sparsity of frame-differences for efficient measurements in a tracking system has been recently suggested and demonstrated in a linear imaging model [19]. Our current work deals with a different group of problems, namely, is a quadratic mathematical problem (phase retrieval). For such nonlinear problems, standard linear sparse recovery algorithms are inapplicable. This, together with the motivation to pursue sub-wavelength imaging, are the reasons we develop the new methodology presented here. Applications of our method may include imaging of moving or changing biological specimen, magnetic nano-structure formation, or multiple particle tracking, in all of which the improvement of temporal resolution can be greatly beneficial. Imaging magnetic domains during their formation is particularly appropriate for our method because the temporal variations are sparse, as the domains tend to flip in small groups. Another interesting application where our method can be used is for CDI of growing neurons [20] – where, due to the nature of the dendritic growth in which most of the neuron remains unchanged, the difference between consecutive frames is sparse. Yet another example for a phenomenon that conforms to our sparse-variation model is the formation of crystal, such as the colloidal crystal dynamically imaged by STED [21]. Our method relies on new ideas associated with sparsity-based subwavelength imaging [22–24] and phase retrieval of sparse objects [25–30], and utilizes a recently developed efficient greedy algorithm [30] in order to solve the problem.

2. CDI problem formulation

In plane-wave CDI, an object is illuminated by a plane wave, and its far field diffraction pattern is measured, as depicted in Fig. 1. Therefore, performing CDI amounts to recovering the object's optical transmission function from the magnitude of the diffraction pattern. In the far-field, or the Fraunhofer approximation (applicable when $a^2/\lambda d \gg 1$ and $b^2/\lambda d \gg 1$, d being the distance between the detection plane and the object, and a and b being the radii of circles within which the object and its far-field diffraction pattern are localized around the optical axis, respectively), the intensity of the diffraction pattern corresponds to the magnitude of the Fourier transform of the object [31]. Such is the case when the features of the object are large compared to λ, and the diffraction pattern is measured sufficiently far away.
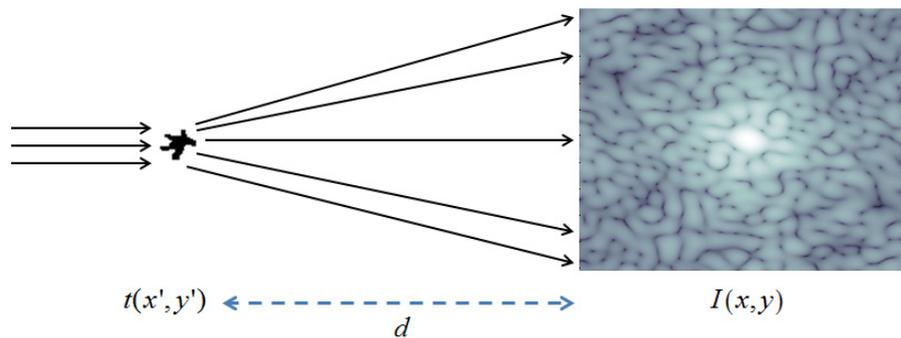

Fig. 1: CDI setup

Denoting the 2D transmission function of the object as $t(x',y')$ and the measured intensity as $I(x,y)$, under the assumptions above, the following relation holds, up to a constant multiplicative factor:

$$I(x,y) = \left| T\left(\nu_x = \frac{x}{\lambda d}, \nu_y = \frac{y}{\lambda d} \right) \right|^2 \tag{1}$$

where $T(\nu_x, \nu_y) \triangleq \iint t(x',y') \exp\left[j2\pi(\nu_x x' + \nu_y y') \right] dx' dy'$ is the 2D Fourier transform of $t(x',y')$. The mathematical description of the CDI problem is therefore to invert Eq. (1), namely to recover the phase of $T(\nu_x, \nu_y)$ - since once its phase is known, it can be inverse-Fourier transformed to yield $t(x',y')$. However, due to the missing phase information, the inversion problem is ill-posed, and in general some additional prior information is necessary in order to solve it. Most commonly, this additional knowledge is about the location of the object and its boundaries (its "support"). The support is incorporated into an iterative phase retrieval algorithm developed by Fienup [7], attempting to find a solution consistent with both the prior information and the measurements. There are several variants of the Fienup algorithm [8], and they all work by alternating between imposing real-space constraints (e.g. the assumed support of the object) and Fourier constraints – i.e. the measured Fourier magnitude. In order to converge to the correct solution, however, they require both oversampling, i.e. a larger number of measurements than unknowns (more than twice), and high SNR [10].

3. Sparsity-based dynamic CDI

Using CDI to image a dynamic process requires taking a set of far-field measurements at consecutive times. In principle, in each temporal segment, the far-field diffraction pattern can be captured, its phase can be retrieved using a phase-retrieval algorithm [7], and a snapshot of the object can be recovered. However, due to the high SNR requirement of current phase-retrieval methods, and in order to avoid possible damage to the object, a strong enough signal can only be obtained by accumulating many weak pulses. The acquisition time and hence the temporal resolution are therefore inherently very slow – on the order of tens of seconds at least [14,17,18], rendering current CDI techniques inapplicable for fast imaging dynamic processes.

Here, we suggest and demonstrate a method facilitating considerably shorter acquisition times for CDI of dynamically varying objects, and in doing that allowing for improving the temporal resolution substantially. Our technique works under the assumption that two consecutive frames are similar, rather than treating each temporal frame separately. Under this assumption, which is commonly used in applications of video compression [32] and tracking systems [19,33], the number of unknowns becomes smaller than the number of measurements. At the k'th frame we have the following system of quadratic equations:

$$I^k(x,y) = \left| \mathsf{F}\left[t^k(x,y) \right] \right|^2 = \left| \mathsf{F}\left[t^{k-1}(x,y) + \Delta t^k(x,y) \right] \right|^2 \tag{2}$$

where $I^k(x,y)$ is the k'th temporal frame, and F denotes the 2D Fourier transform. The difference between consecutive frames, denoted as $\Delta t^k(x,y)$, is therefore sparse, i.e. it contains a small number of nonzero elements, when it is spatially discretized.

In our method for solving the dynamic CDI problem, we shall assume that the first frame is measured with high quality – i.e. high SNR and sufficient number of samples. However, from then on the following frames can be sampled under less strict conditions (low SNR, fewer samples). Under this assumption, after spatial discretization, Eq. (2) translates to finding a sparse solution to a set of quadratic equations.

Finding sparse solutions to underdetermined systems of equations is a problem that has been studied extensively in recent years - although until recently all such studies addressed linear equations only (see, e.g., [34–37]). In fact, the first studies on finding sparse solutions to sets of quadratic equations appeared only in the past two years, in the context of sparsity-based sub-wavelength imaging [38,39], and have recently been considered in more general

contexts [25,26]. In the context of optical imaging, sparsity has also been used in various different application such as digital holography [40,41], phase space tomography [42], tracking [19,33], and more. Here, in order to solve the quadratic problem, we use a recently developed greedy method, based on iterative application of the damped Gauss Newton algorithm [30].

4. Formulation of the discrete problem

Discretization of Eq. (2) yields the following vector relation:

$$y_k^i = \left| F^i x_k \right|^2 = x_k^* F^{i*} F^i x_k = \left(x_{k-1} + \Delta x_k \right)^* F^{i*} F^i \left(x_{k-1} + \Delta x_k \right), \qquad (3)$$

where x_k is a $n^2 \times 1$ vector representing the $n \times n$ sought object in time k, obtained from the discretized 2D object by stacking its columns, and the vector y_k is, similarly, a $m^2 \times 1$ column stack of the Fourier intensity measurements. The notation F^i is the i'th row of the 2D DFT matrix, so that each element i in y_k is obtained by the absolute value squared of the inner product of x_k with the appropriate spatial frequency. Finally, Δx_k is the frame difference vector, i.e. $\Delta x_k \triangleq x_k - x_{k-1}$. Equation (3) can be rewritten as a function of the unknown frame-differences:

$$y_k^i = x_{k-1}^* A^i x_{k-1} + \Delta x_k^* A^i x_{k-1} + x_{k-1}^* A^i \Delta x_k + \Delta x_k^* A^i \Delta x_k,$$
$$(4)$$

where A^i is a matrix defined by $A^i \triangleq F^{i*} F^i$. The formulation of Eq. (4) shows that, given the recovery of the previous frame, x_{k-1}, the measurements are a quadratic function of the frame-difference vector Δx_k only. The goal is therefore to find a solution Δx_k that is consistent with the measurements. However, in general, this problem is ill-posed, and the solution Δx_k is not unique. We therefore regularize the problem by requiring the frame-difference to be sparse, based on the logic that the object changes only slightly between consecutive time segments. As we will see, this essential concept of sparsity enables improved performance over standard methods.

The mathematical formulation we propose is summarized as follows. At each time-segment k, we seek a vector Δx_k that is consistent with the quadratic measurement vector y_k (through the relation given by Eq. (4)), and is at the same time sparse – i.e. contains a small number of non-zero elements. This yields the following optimization problem:

$$\widehat{\Delta x}_k = \arg\min_{\Delta x_k} \sum_{i=1}^m \left(y_k^i - \left(x_{k-1} + \Delta x_k \right)^* A_i \left(x_{k-1} + \Delta x_k \right) \right)^2$$
$$\text{subject to } \left\| \Delta x \right\|_0 \leq s \qquad (5)$$

where $\left\| \cdot \right\|_0$ denotes the zero-norm, i.e. the number of non-zero elements, and s is a parameter representing an upper limit on the number of nonzero elements in Δx.

5. Solution method – GESPAR

We find a solution to Eq. (5) by using a slightly modified version of our recently developed greedy algorithm for sparse phase-retrieval, named GESPAR, described in detail in [30]. The algorithm is based on a local search, where within each iteration we solve a nonlinear least-squares problem using the damped-Gauss-Newton (DGN) [43] method. Given the previous frame x_{k-1}, the algorithm produces a solution Δx that is both sparse and consistent with the quadratic measurements of the current frame. The algorithm's steps are summarized below:

1. Select an initial random support for Δx_k, namely the initial locations of the nonzero elements in the guessed solution.

2. Given the support, the problem defined by Eq. (5) reduces to a nonlinear least squares problem, which we solve by the DGN algorithm that is commonly used for this type of problem. The DGN produces an estimate Δx.

3. Calculate the cost function $f(\Delta x) = \sum_{i=1}^{m} \left(y_k^i - (x_{k-1} + \Delta x)^* A_i (x_{k-1} + \Delta x) \right)^2$, along with its gradient around the current estimate Δx.

4. Perform a local search by index swapping: replace an index i from the support of Δx containing an element with a small absolute value element with an index j containing a high absolute gradient value. Given the new support s, perform DGN, obtain Δx_s, and calculate the cost function $f(\Delta x_s)$.

5. Update support: If the cost function is improved, i.e. $f(\Delta x_s) < f(\Delta x)$, then the support is updated - i.e. the index i is removed from the support and the index j is added to it . Go to step 3. Otherwise, go to 4, until a predetermined number of swaps have been performed. If there is no possible improvement in the cost function by swapping - go to step 1.

The procedure stops after a predetermined number of swaps. The output is the solution Δx that yields the lowest value of the cost function $f(\Delta x)$. Since the algorithm recovers the frame-difference at each temporal segment, we avoid carrying a large accumulative error by assuming that the initial frame x_0 is recovered with high quality, i.e. using sufficient measurements and high SNR. The main difference between the method used here and GESPAR as detailed in [30] is the fact that here the cost function in the current problem contains an additional term linear in Δx.

It is important to note that problem (5) is very difficult to solve - it is highly non-convex and has combinatorial complexity. In general, there is no guarantee that the true solution or a global minimum will be found, and indeed GESPAR finds only a local minimum. However, when the sought signal is sparse enough, empirical evidence (based on numerical simulations in [30]) suggests that the true solution is found with high probability. In order to avoid stagnation of the algorithm, for reasons such as the "twin image problem" [44,45] associated with iterated projection algorithms [46], GESPAR includes randomization. Specifically, the indices i and j from part 4 in the algorithm are selected randomly from a sub-set of indices.

6. Simulations

We test the performance of our algorithm by performing a series of numerical simulations. The first simulation examines robustness to noise: The Fourier transform magnitude of a simulated object that is varying in time is sampled under various SNR conditions. The real-space object is then recovered frame by frame, and the recovery is compared with the true dynamic object. The recovery result is compared with frame-by-frame recovery using the

standard Fienup Hybrid Input-Output (HIO) algorithm [8]. The size of the simulated object is 51x51 pixels, and the sparsity of each frame-difference is $s = 5$. The number of Fourier measurements taken for each frame in this example is 2601. In all of these simulations the number of iterations used in the Fienup HIO algorithm is chosen as 3000, which is sufficient for convergence, where the last 600 iterations are averaged to yield the final reconstruction. The parameter β in the HIO algorithm is chosen as 0.5, which yielded the best results. The initial guess for HIO in all frames is the clean first frame, which proves to work better than other options, such as taking the initial guess as the output of the previous frame. The prior support knowledge, needed for the HIO algorithm, is taken as the 31x31 pixel square, visible in the videos and in Figs. 2 and 4. The total number of GESPAR swaps is set as 500 for each frame.

Simulated examples showing our sparsity-based reconstruction of dynamically varying objects, and comparing to standard HIO reconstruction are shown in Fig. 2 (Media 1) and Fig. 2 (Media 2). Figure 2 shows two example frames: the left panel shows the true simulated object, the center shows our GESPAR recovery, based on frame-difference sparsity, and the right panel displays the HIO recovery. In this example, the SNR is 30dB, defined as: $SNR = 20\log\left[\|y\|_2 / \|n\|_2\right]$, where n is an independently and identically distributed (i.i.d.) Gaussian noise added to the measurement vector y. In these typical examples GESPAR clearly obtains a much better reconstruction than the HIO algorithm, as will be quantified next.
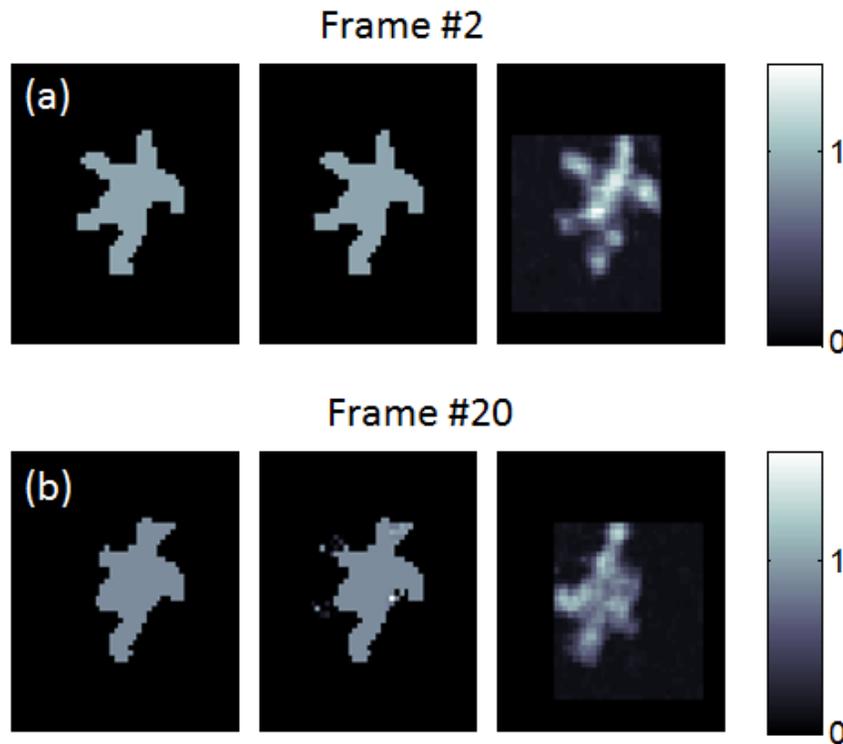

Fig. 2: Simulated example frames comparing the quality recovered through our sparsity-based phase-retrieval technique with the best result calculated through the Fienup HIO algorithm. (a) Frame #2, and (b) Frame #20. The left column shows the true simulated frames, the center frame shows our GESPAR reconstruction (Media 1), and the right column shows the reconstruction with the standard Fienup HIO (Media 2). SNR=30. The simulated object has uniform intensity and phase – i.e. it is opaque. However, the asymmetry of the object means that its phase structure in Fourier space is complicated and not binary.

The performance of both algorithms under various SNR is displayed in Fig. 3. The normalized mean squared error, defined as $NMSE = \left\| x_{rec} - x \right\|_2 / \left\| x \right\|_2$, where x is the unknown object and x_{rec} is the reconstruction, is plotted as a function of frame number for different SNR's.
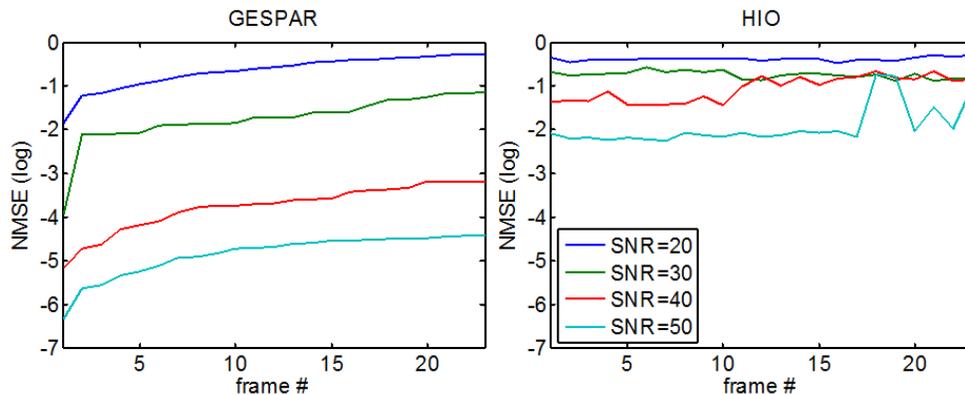

Fig. 3: Reconstruction error (log NMSE) for various noise values for our GESPAR (left) algorithm and the Fienup-HIO (right)

Several conclusions can be drawn from Fig. 3. First, in order to obtain a solution of similar quality (measured by NMSE), the HIO algorithm requires a much higher SNR than GESPAR algorithm; e.g. the performance of GESPAR under SNR of ~30 corresponds to SNR of ~50 in HIO, corresponding to a factor of 10 in the relative signal to noise energy. Since the SNR is related to acquisition time, this translates directly to reduction in acquisition time. Namely, under the reasonable assumption that the dominant noise factor in the imaging system is shot-noise, the factor of 10 in signal to noise energy translates to a factor of 100 in acquisition time. Second, the visible increase in the error with frame number can be explained by the fact that GESPAR iteratively reconstructs only the differences between frames, so the error is accumulated. The less apparent increase in the HIO error with frame number is the result of the initial guess being more similar to the first frames than to the following frames in our example.

In addition to the considerably superior performance of our algorithm in terms of noise robustness, we also examine the performance as a function of the number of measurements required for reliable reconstruction. Once again, successful recovery from a small number of measurements directly leads to faster acquisition time. For example, several adjacent measured pixels can be summed ("binned") together, yielding a smaller number of measurements but a higher SNR. Figure 4 shows two example frames recovered using only 650 Fourier magnitude measurements (about ¼ of the number used in Fig. 2), located on a square grid, and with very low noise (SNR=60). Fienup algorithms are known to require over-sampling in the Fourier plane [10], therefore the low quality of the HIO reconstruction in Fig. 4 (Media 4) (right column) is not surprising. Using the sparsity prior information, on the other hand, compensates for the considerably lower number of Fourier measurements, leading to the high quality reconstruction with GESPAR shown in Fig. 4 (Media 3) (center column). Figure 5 shows the reconstruction error more quantitatively, as a function of frame number, for both algorithms. The reconstruction error increases as a function of frame number in both algorithms, for related but different reasons. In the HIO algorithm the reason is that the initial guess used for each frame is the clean first frame – which proved to work better than using a random initial guess or the output of the previous frame. With time, the object changes more and more, so that the first frame becomes a worse and worse initial guess. In contrast, for GESPAR, the increase in reconstruction error is mainly due to the accumulating error, which results from reconstructing only frame-differences, as discussed earlier.
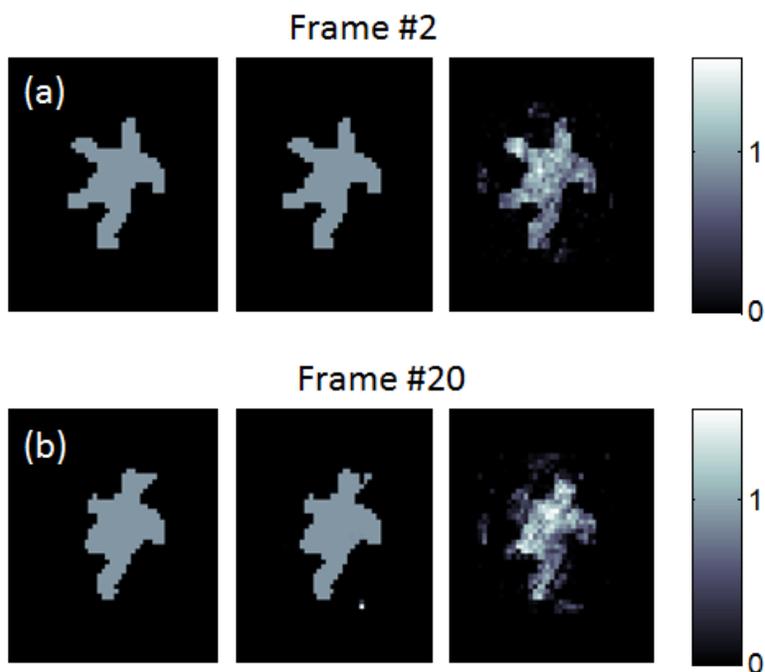
Fig. 4: Simulated example frames comparing the recovery using 650 measurements through our sparsity-based phase-retrieval technique with the best result calculated through the Fienup HIO algorithm. The left column shows the true simulated frames, the center frame shows our GESPAR reconstruction (Media 3), and the right column displays the recovery using the standard Fienup HIO reconstruction (Media 4). SNR=60. The simulated object has uniform intensity and phase – i.e. it is opaque.
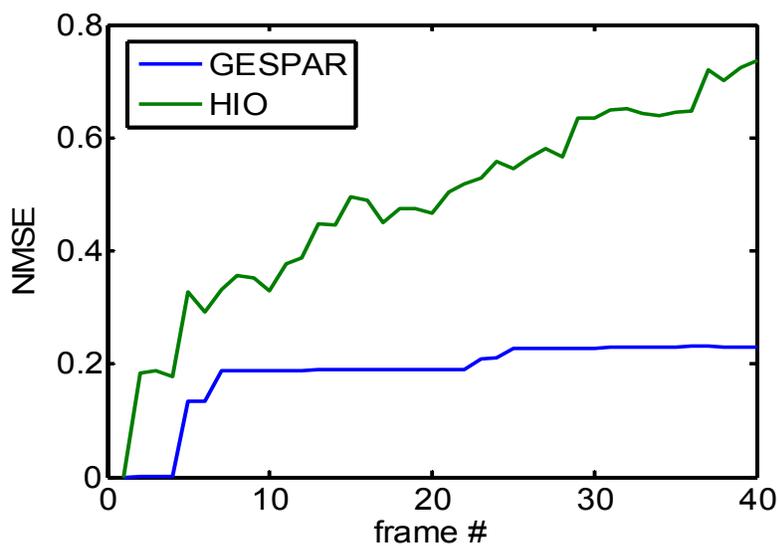

Fig. 5: Reconstruction error vs. frame number using GESPAR and HIO algorithms. The number of samples is 650. SNR=60.

We also test the performance of our GESPAR algorithm under various levels of sparsity, as shown in Fig. 6. The reconstruction quality naturally improves as the frame difference is more sparse (low values of s).
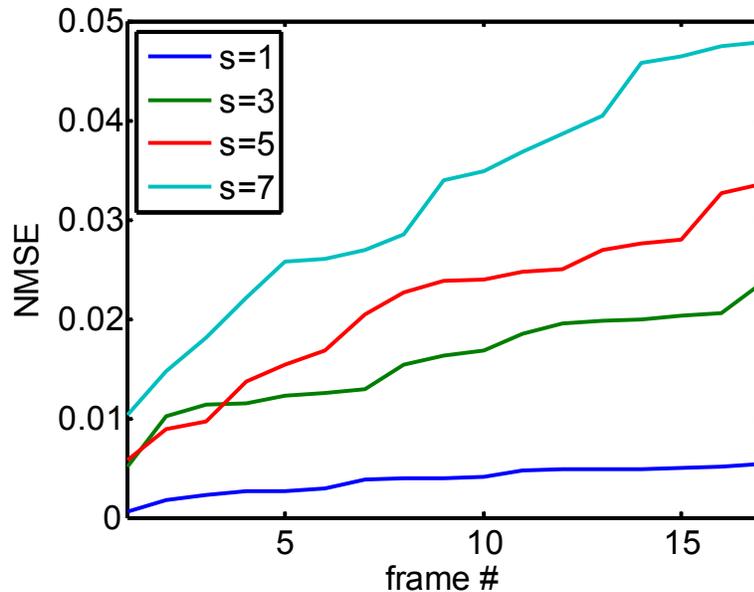

Fig. 6: Reconstruction error vs. frame number for various sparsity levels. SNR=40. Error accumulation due to the reconstruction of frame-differences can impose a practical limit on the possible number of frames in a reconstructed movie. This can be tackled by performing high-quality sampling every certain number of frames, dictated by the system's parameters (e.g. noise, number of measurements). In this case, it is also possible to further improve performance by back-propagation - i.e. using a high-quality frame in order to reconstruct a previous frame.

The spatial resolution of our technique is an important issue. One may naively think that the best resolution is pre-determined by the high-resolution first frame. However, our recent work on sub-wavelength imaging demonstrates that it is possible to obtain resolution beyond the resolution of the measurement device, which in our current case would correspond to improving the resolution beyond the resolution of the first frame. Indeed, current work is targeting this goal.

The limitations of our technique may arise from several causes. Of course, in some cases the assumption that the frame-difference is sparse may be inappropriate. Such would be the case, for example, if the images of the dynamically changing object are taken at a low rate – allowing the object to change significantly between two consecutive frames. Another limitation may arise from the fact that the GESPAR algorithm recovers only frame-differences – implying that the high-resolution first frame is important, and that error will generally accumulate with frame number – as our simulation results indeed indicate. The temporal resolution of the imaging system, although improved significantly by our method, will still ultimately depend on SNR.

7. Summary and future directions

We suggested and demonstrated a method to perform fast dynamic CDI, based on the prior knowledge that the object changes only slightly (sparsely) in between successive frames. The problem was formulated as seeking a sparse solution to a set of quadratic equations, and a recently developed algorithm for sparse phase retrieval (GESPAR) is used to solve it. In principle, additional prior knowledge about the object can be incorporated into the method in

order to further improve the results – with the ultimate goal of relieving the method of the need for a clean first frame.

Importantly, since our method succeeds in recovering the features of an object from only a subset of its Fourier spectrum, it can serve as a solid basis for the recovery of sub-wavelength features. This is because the problem of sub-wavelength CDI corresponds to recovering an object using only the magnitude of its low Fourier frequencies [39]. Indeed, our next goal is to employ the current ideas to obtain super-resolution beyond the numerical aperture of the system.

Acknowledgments