SIAG/OPT Views and News

A Forum for the SIAM Activity Group on Optimization

Volume 25 Number 1 January 2017

Contents

Comments from the Editors

We thank the contributors to the 35th issue of SIAM Activity Group on Optimization's newsletter. The theme of this issue is "optimization and imaging", with an article on phase retrieval by D. Russell Luke and nonnegative matrix factorization by Nicolas Gillis. We hope you enjoy this latest installment of Views and News.

Congratulations to the new officers for the SIAM Activity Group on Optimization: Tamás Terlaky, Andreas Waechter, Michael Friedlander, and James Luedtke. Thank you for agreeing to serve our activity group for the next three years.

Many of you have written to opt for an electronic copy of Views and News; for the others among you, please do not hesitate to contact us to opt out of receiving physical copies.

As always, we welcome your feedback, (e-)mailed directly to us or to siagoptnews@lists.mcs.anl.gov. Suggestions for new issues, comments, and papers are always welcome!

Stefan Wild, Editor
Mathematics and Computer Science Division, Argonne National Laboratory, USA, wild@anl.gov, http://www.mcs.anl.gov/~wild
Jennifer Erway, Editor
Department of Mathematics, Wake Forest University, USA, erwayjb@wfu.edu, http://www.wfu.edu/~erwayjb

Articles

Phase Retrieval, What's New?

D. Russell Luke
Institut für Numerische und Angewandte Mathematik,
Universität Göttingen, 37083 Göttingen, Germany.
E-Mail: r.luke@math.uni-goettingen.de
http://num.math.uni-goettingen.de/~r.luke

Ask an engineer to solve a problem and she will come back in a day or so with something that seems to work well enough most of the time. Ask a mathematician to solve the same problem and he will return many months later with an exact but unimplementable solution to a different problem. I'm sure most readers of this newsletter have heard some variation of that joke. But a true story lies somewhere in there, a story that is writ large with the phase retrieval problem.

The phase retrieval problem has been around for more than a century, and it is solved tens of thousands of times each second, mostly by physicists. Phase retrieval plays a central role in the x-ray imaging experiments conducted by researchers here in Göttingen, where we are in the last 5-year funding cycle of a 15-year collaborative research center studying nanoscale photonic imaging (Deutsche Forschungsgemeinschaft CRC755). The center consists of experimental physicists and biomolecular physicists building new instruments and observation techniques (one of those techniques, STED, won center participant Stefan Hell a Nobel Prize in 2014) as well as mathematicians studying algorithms, image processing, and statistics. Phase retrieval is an applied mathematician's dream problem: it is central to many imaging modalities, it is simple to state, numerical routines for its solution abound, and it is mathematically interesting in ways that solving systems of linear equations will never be.

Nick Trefethan wrote in his introduction to a 2002 SIAM Review article on phase retrieval that I wrote together with Jim Burke and Rick Lyon [42], "A Google search of 'phase retrieval' returns 271,000 records." Almost fifteen years later, a Google search returns 364,000 records (with the safe search on). A Web of Science™ database search back to 1945 yields 8,924 results, 7,189 of those since 2002, more than half of those since 2011. A lot of new interest has been expressed recently in particular within some corners of the statistics and applied mathematics communities. Apparently, money also is at stake, as I learned in preparation for this article. So, if the problem has been around for so long and people are already solving it–Hauptman and Karle won a Nobel Prize in 1985 for solving the crystallographic phase retrieval problem–what's all the recent fuss about?

For those who don't know what phase retrieval is, it is simply stated as follows: Find $x \in C \subset \mathbb{C}^n$ such that $|(F_k x)|_j = b_{jk}$, where for $k = 1, 2, \ldots, K$ the mapping $F_k : \mathbb{C}^n \to \mathbb{C}^m$ is linear and $b_{jk} \in \mathbb{R}_+$ for all $j = 1, 2, \ldots, m$ and $k = 1, 2, \ldots, K$. The classical problem comes from diffraction imaging where the set C is some a priori constraint like support or nonnegativity, and the mapping F_k is a Fourier transform of some kind. This includes the Fresnel transform and some defocused or otherwise imperfect Fraunhofer transform. The structure of the model is simple: the components of the vectors $F_k x$ must lie on circles of a given radius in the complex plane. Unfortunately, circles are not convex (the sphere does not contain any line segment joining any two points on the sphere), causing all sorts of problems, both mathematical and practical. The first statement of the phase problem that I could find goes back to a letter from J. W. Strutt (Lord Rayleigh) to Michaelson in 1892 [53]. Lord Rayleigh was pessimistic about the prospects of breaking the phase barrier unless a priori information about symmetry was known. A solution to the band-limited phase problem is a zero of a related complex polynomial. I'll come back to the qualifier "band limited" in a moment, but ignore that detail for now. In the 1950s Akutowicz [1, 2] showed that the one-dimensional phase retrieval problem without any a priori constraints has many solutions, since by the fundamental theorem of algebra, all 1D complex polynomials factor into products of monomials. A lot of workarounds for the unconstrained 1D phase retrieval problem have been developed since the 1970s [17, 25, 31, 50, 52, 57, 58, 7], all of which involve adding a constraint implicitly or explicitly. Some recent progress on the 1D problem has come from initialization techniques that land one in a neighborhood of the minimum phase solution [30] where the usual nonlinear programming techniques can perform reliably.

At the end of the 1970s Bruck and Sodin [10] pointed out that the fundamental barrier to unconstrained 1D phase retrieval does not apply in higher dimensions since, magically, polynomials of dimension two or more *almost never* factor. This conformed nicely with the unreasonable success of the simple Gerchberg-Saxton [23] and HIO [21] algorithms for 2D phase retrieval proposed a few years earlier. Shortly thereafter Hayes [24] proved that for band-limited signals, the 2D phase retrieval problem has unique solutions, almost surely, up to rotations, shifts, and reflections. One might conclude that the book on phase retrieval was closed a long time ago, except that the theory didn't quite match up with practice as nicely as one would hope. The first hint that the story is more complicated came from the algorithms themselves. They worked fairly well a lot of the time, but one of the more popular approaches, HIO, *never* worked in the usual sense of convergence to a fixed point. Still today, people continue to apply HIO according to the following recipe: run 10-40 iterations of HIO; then apply several passes of Gerchberg-Saxton to clean up the image; publish.

I have the fortune of being coauthor with Heinz Bauschke and Patrick Combettes of a paper on phase retrieval algorithms that gets a steady stream of citations [3]. Unfortunately, fewer people read it than cite it. We started from the premise that really only a handful of good first-order algorithms exist and that anything that works is probably a tweak of one of those. It was known before our paper that Gerchberg-Saxton and the *error reduction algorithm* [21] are simply alternating projections (one of the handful of good algorithms). We were able partially to identify HIO for the case of a support constraint alone by showing the correspondence between this procedure and the now ubiquitous Douglas-Rachford algorithm. In a follow-up paper [4] we showed that HIO with a support and nonnegativity constraint becomes a *different* fixed-point iteration, what we called the hybrid projection reflection (HPR) method (*not* one of the handful). This fundamental change in the fixed-point mapping by a seemingly minor change in the constraint structure is not obvious when the algorithms are written in the format favored in the optics literature. At the same time the HPR method was presented, Veit Elser introduced his *difference map* [19], which for certain parameter values coincides with the Douglas-Rachford and HPR algorithms, again depending on the constraint structure [38]. The instabilities of these algorithms together with the insight provided by the more mathematical prescription of the algorithms led me to propose a relaxation of the Douglas-Rachford algorithm, which I called RAAR, that has fixed points when Douglas-Rachford does not [39]. At that time, alternating projections, Douglas-Rachford, HPR, and RAAR algorithms were understood only for convex problems. In the convex setting Douglas-Rachford can be identified with the alternating directions method of multipliers [22], which is currently popular for large-scale problems. For nonconvex problems, however, our understanding of Douglas-Rachford and even alternating projections, and hence everything else close to these, pretty much evaporated. Since then a lot of quiet, patient work has been done in the variational analysis community to develop the theory of first-order methods for nonconvex problems, and much of the missing theory behind the success of these algorithms for phase retrieval is in place. But I get ahead of myself.

Almost any article on phase retrieval in the applied mathematics literature will start with a statement like "The phase retrieval problem is found in many different areas of science and engineering, such as x-ray crystallography, astronomy, diffraction imaging, and more." So I contacted several physicists and astronomers to find out from them what is new in phase retrieval. One place where efficient solutions to the phase retrieval problem is of vital importance is the W. M. Keck Observatory. The Keck instruments need to correct for random aberrations in the Earth's atmosphere in order to compete with instruments such as the Hubble Space Telescope. The shape of the atmosphere is encoded in the phase of the observations. Sam Ragland, an adaptive optics scientist at Keck, told me that their instruments use Shack-Hartmann sensors to measure the phase directly. This is a hardware solution to the phase problem. Keck's wavefront controllers operate at a rate of 2 kHz, 1 kHz in practice. Ragland did say that they were testing a computational phase-diversity algorithm (phase retrieval with several defocused images), for which they are at the moment using just the Gerchberg-Saxton algorithm.

The cost of the Shack-Hartmann sensors is photons, which are in short supply in more modern x-ray imaging set-ups. The dominant approach here is computational phase retrieval, although recent proposals involve adding random masks to the imaging systems in order to avoid nonuniqueness in the numerical reconstructions. With regard to algorithms for phase retrieval, Elser, an expert in diffraction x-ray imaging at Cornell University, has seen "nothing significant" in the past 5 years. He is even more dismissive of the impact of random masks: "The integrity of those phase masks has to be established at the same resolution as their intended application – and for that you need a phasing system that works at that resolution!" This is a fundamental challenge for nanoscale (i.e., x-ray) imaging. Pierre Thibault, an expert in blind ptychography (the second worst phase retrieval problem one can encounter) at the University of Southampton, is more circumspect. Algorithms such as alternating projections, HPR, Douglas-Rachford, RAAR, and the difference map work just fine, according to Thibault; "the biggest bottleneck is hardware." While Thibault has found no use for random masks in his work, he would not discount the possibility that the idea of randomly generated data could have unforeseen applications. One such possibility he mentioned was Fourier ptychographic microscopy, which iteratively stitches together a number of variably (i.e., randomly) illuminated, low-resolution intensity images in Fourier space to produce a wide-field, high-resolution complex sample image [56, 59]. The idea that more Fourier measurements can improve the performance of phase retrieval algorithms has been around for a while. For crystallographic phase retrieval there is only so much information you can get out of the intensity measurements. But in the early 2000s it was recognized that for noncrystallographic measurements one is not limited to sampling on a fixed lattice and that oversampling dramatically improves numerical reconstructions [45]. Unfortunately, this improved performance is attributed to some form of uniqueness. This is curious since a few moments of reflection on elementary Fourier analysis is all that is needed to be convinced that oversampling has nothing to do with uniqueness. Increased, but still finite, sampling in the Fourier domain just pushes the error created by trying to reconstruct a compactly supported object from finitely many Fourier measurements to some level below either numerical or experimental precision.

The oversampling justification is just one example of a fixation on uniqueness that has overshadowed the most obvious structural problem for phase retrieval: existence. Remember the "band-limited" qualifier in Hayes' result cited earlier. What that means is a compactly supported Fourier transform. And, what that means is that the object must be periodic. At this point many people retreat to the discrete Fourier transform, which is a unitary linear operator, about the best kind of operator one can have – except that the best physical model we have for describing what we measure in any physical experiment is a sample of the continuous Fourier transform. And when you implement phase retrieval on a computer, you cannot avoid implicitly setting the values of the part of the Fourier spectrum that you do not measure to zero. So almost any constraint–in particular compact support–that you place on the object whose Fourier transform you sample will be inconsistent with the measurement. One exception is crystallography, where (perfect) crystals are indeed periodic and the Fourier transform can be assumed reasonably to be band limited. The most exciting and difficult imaging challenges today, however, are in noncrystallographic "single-shot" x-ray imaging [48, 20]. I mentioned that blind ptychography is the second hardest type of phase retrieval problem you might encounter. Blind ptychography [26, 55, 44, 29] is akin to reading a fragment of an ancient text in a script you have never seen, with a pair of glasses borrowed from someone you have never met, and being asked to reconstruct simultaneously the script and the prescription of the glasses. For single-shot x-ray imaging the script consists of 3D figurines floating randomly in midair and of which you get only brief flashes from a strobe light. Phase retrieval is the easy part for reconstructions from single-shot data - the challenging part is the tomographic reconstruction of the Fourier data. Unfortunately, for much of the more recent mathematical work directed at phase retrieval to have any traction, uniqueness is essential; but for these more modern applications even existence of a solution to the model equations together with qualitative constraints cannot be taken for granted.

The recent work in applied mathematics on phase retrieval has its roots in a series of now-famous papers by Candès and Tao [14, 13], which showed that under certain conditions on the matrix generating an affine subspace of $\mathbb{R}^n$ (called the *restricted isometry property* in the literature), there is a unique sparsest point in the subspace and, moreover, this point is the point with smallest ℓ_1-norm. When the space is a space of matrices, the uniqueness is up to orbits, and the elements of the orbits have smallest nuclear norm. This has sparked a wave of papers on convex (and even nonconvex) relaxation in the signal processing community since finding a point in an affine subspace with minimum norm is a convex optimization problem while the problem of minimizing the counting function subject to an affine constraint is nonconvex and NP-hard [47]. Bloomensath and Davies [8, 9] ran against the current, however, and examined a simple forward-backward prox-algorithm, iterative hard thresholding, for solving a slightly different problem of minimizing the norm of the residual in the image space subject to a sparsity inequality constraint in the domain. They showed that under an asymmetric generalization of the restricted isometry conditions required for the correspondence of the nonconvex sparsity optimization problem and its convex relaxation, iterative hard thresholding converges globally to the global solution of the sparsity-constrained residual minimization problem. Their work inspired me and my students Robert Hesse and Patrick Neumann to show that the same, or similar, conditions also guarantee global linear convergence of alternating projections for the problem of finding the intersection of an affine subspace and the set of vectors with sparsity less than a given value [28]. We also showed that the asymmetric restricted isometry conditions for global convergence of alternating projections imply *transversality* of the range of the transpose of the matrix generating the affine subspace and the orthogonal complement of all subspaces of dimension twice the dimension of the sparsest element. A consequence, which has not been much explored, is that there cannot be any locally optimal solutions for this nonconvex problem other than the unique global minimum. Apparently, convex relaxations are not even needed for problems with this structure.

The connection to phase retrieval, pointed out first by Candés, Eldar, Strohmer, and Voroninski [12], uses a well-known trick from conic programming for turning a quadratic function on $\mathbb{R}^n$ into a linear function on the space of $n \times n$ matrices. The price to pay for this is in going from a problem with n unknowns to a problem with n^2 unknowns. Structurally, however, the problem in the space of matrices is the same as the spasity optimization problem above. There was some hope that this would lead to a breakthrough in phase retrieval by solving a convex problem in the space of matrices. But the prospect of squaring the number of unknowns should have given pause to even the most ardent booster of Moore's law. Phase lift, as this idea is called in the literature, has not proven to be a reasonable computational strategy. In the past two years there has been some backing away from phase lift as an algorithm, and more direct nonconvex methods are again being proposed (see [32] and references therein). No one from the applications side that I spoke to for this article was aware of these newer methods, however. One reason could be that none of the newer methods has been compared with the methods favored in the optics community. Reference to methods such as Gerchberg-Saxton and HIO and the identification with classical algorithms made in [3] seems to be obligatory in recent articles, but they appeal to missing theory behind these methods in to order avoid direct comparisons.

I'm happy to report that we actually now know quite a bit about the classical algorithms for the phase retrieval problem. I mentioned in the beginning that only a handful of good first-order algorithms exist. Among these are steepest descent (which includes averaged projections, the Misell algorithm for phase diversity, and many of the schemes proposed in the past few years), forward-backward prox schemes (which include projected gradients, hard- and soft-thresholding, and accelerations), backward-backward prox schemes (which include alternating projections and hence Gerchberg-Saxton), and Douglas-Rachford (to which category I assign HIO, RAAR, and the difference map). The standard "old" algorithms for phase retrieval are all based on projectors that are composed and averaged in some fashion. Implicit in this is a *feasibility* formulation of the phase problem, that is, to find some point in the intersection of the set of points satisfying the constraints implied by the data measurements and the set of points satisfying qualitative constraints such as support and nonnegativity. This is an extremely powerful modeling approach since it is easy to introduce new constraints without changing one's algorithmic approach. It also lays bare the success and failure of various methods.

When one settles on a feasibility model for a problem, uniqueness is almost irrelevant. The two important cases for feasibility are consistent and inconsistent. In the consistent case the sets have at least one point in common; in the inconsistent case, the sets have no point in common. Most of the progress on the nonconvex convergence theory for the good algorithms above has been for the consistent case. Based on a series of papers exposing the structure of alternating projections in increasingly inhospitable environments [15, 35, 34, 40, 5, 18, 49, 27], we know now that alternating projections applied to consistent phase retrieval problems is locally linearly convergent at points of intersection except in the unlikely case that the constraints are tangential. The nonconvex theory of the Douglas-Rachford algorithm for consistent problems is also fairly well understood in settings that cover phase retrieval [27, 51].

As I argued above, however, the real-world phase retrieval problem is hopelessly inconsistent. This inconsistency can be observed by simply running the Douglas-Rachford algorithm on your favorite experimental data set. You will observe the iterates moving around seemingly chaotically, sometimes looking like something well structured before wandering off to nonsense. This behavior is a consequence of the fact that the Douglas-Rachford mapping applied to inconsistent feasibility has no fixed points. For convex problems this is not a serious issue since the *shadow sequence* of the iterates can be shown to converge [6]. For nonconvex problems, however, all bets are off, and this explains the instability of the HIO algorithm for phase retrieval.

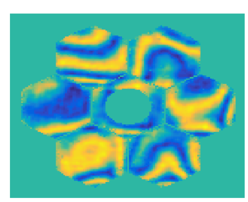

Figure 1: Representative iterate of a noisy JWST test wavefront recovered with the Douglas-Rachford algorithm. For a movie showing instability of the algorithm, go to `http://num.math.uni-goettingen.de/proxtoolbox`

The RAAR algorithm [38, 39] is a relaxation of the Douglas-Rachford algorithm that is guaranteed to have fixed points for a strong enough relaxation. One can easily verify [29, 40] that the regularity of the RAAR mapping for feasibility-based phase retrieval satisfies one of two conditions that together are sufficient to guarantee local linear convergence developed in [51], namely, that the constraint sets are *superregular*. What remains to be shown in order to guarantee that the RAAR algorithm is locally linearly convergent for phase retrieval is that the RAAR fixed-point mapping is *metrically regular* [33] (in some appropriate sense) at fixed points. Abstract formulas for showing metric regularity exist and rely on computing the coderivative of the fixed-point operator [16, 46], but executing these calculations and verifying the conditions for phase retrieval is complicated.

Another way to understand the RAAR algorithm outlined in [39] is as the Douglas-Rachford algorithm applied not to the problem of minimizing the sum of two indicator functions but rather to the problem of minimizing the weighted sum of the squared distance to one of the sets plus the indicator of the other set. Loock and Plonka [36, 37] took this idea further to apply the RAAR algorithm not to the squared distance to a set but rather to the ℓ_1-norm of the wavelet transform of the object to be recovered. They were able to show that the iterates of the RAAR algorithm in this setting are at least bounded. I am optimistic that the remaining open issues concerning the convergence of the RAAR algorithm for phase retrieval will soon be resolved. As pointed out in [39], the power of Douglas-Rachford and its relaxations in the context of feasibility is that it can be tuned to have far fewer fixed points (i.e., locally optimal points), than algorithms such as alternating projections or steepest descents.

For the alternating projections algorithm applied to phase retrieval, the picture is fairly complete. The results are local, as one can expect from any nonconvex problem. For practical, inconsistent phase retrieval, recent work with Matt Tam and Thao Nguyen shows that alternating projections must converge locally linearly to local best approximation points except in rare degenerate cases [43, Theorem 5.10 and Example 5.16]. This result allows for *error bounds* as stopping criteria for this algorithm. What we cannot say is whether the fixed points of the algorithm are good, and this has been the main point of criticism. But let us return to the observation that under conditions similar to, albeit stronger than, those used to justify convex relaxations in sparsity optimization, alternating projections converges globally linearly to a unique global solution in that setting. We can then conjecture that alternating projections for phase retrieval with enough measurements converges globally linearly to a globally optimal best approximation point. A nice opportunity exists here for the two strands of analysis that have been picking away at the phase retrieval problem–one from the variational analysis side and the other from sparsity optimization–to merge productively. A strength of the theory sparked by [14, 13] is that it can say something about how much information is needed before one can reasonably expect nice things to happen on a global scale, and this has nothing to do with convexity or the quantitative local analysis.

The algorithms and phenomena discussed here can be explored in the ProxToolbox [41], which is a slowly growing collection of demonstrations of simple first-order methods built on prox-operators. Following the example of Buckheit and Donoho [11] and more recent calls for reproducible research [54], we are trying to make available all numerical demonstrations that have supported our publications. This effort will expand to data and algorithms from the broader Nanoscale Photonic Imaging Collaborative Research Center at Göttingen. Phase Focus Limited of Sheffield, UK, claims intellectual property rights on iterative routines for ptychography and has sued researchers. No academic researcher has the means to challenge such assertions and this has put a chill on efforts to disseminate information, but it does not appear to be an outright barrier. In this age of increased suspicion of science and the scientific method, it is all the more important to make our work as transparent and accessible as possible.

Acknowledgments. This work was supported in part by German Israeli Foundation Grant G-1253-304.6 and Deutsche Forschungsgemeinschaft Collaborative Research Center SFB755.

REFERENCES

[1] E. J. Akutowicz. On the determination of the phase of a Fourier integral, I. *Trans. Amer. Math. Soc.*, 83:179–192, 1956.

[2] E. J. Akutowicz. On the determination of the phase of a Fourier integral, II. *Proc. Amer. Math. Soc.*, 8:234–238, 1957.

[3] H. H. Bauschke, P. L. Combettes, and D. R. Luke. Phase retrieval, error reduction algorithm and Fienup variants: a view from convex feasibility. *J. Opt. Soc. Amer. A.*, 19(7):1334–45, 2002.

[4] H. H. Bauschke, P. L. Combettes, and D. R. Luke. Finding best approximation pairs relative to two closed convex sets in Hilbert spaces. *J. Approx. Theory*, 127:178–92, 2004.

[5] H. H. Bauschke, D. R. Luke, H. M. Phan, and X. Wang. Restricted normal cones and the method of alternating projections: applications. *Set-Valued Var. Anal.*, 21:475–501, 2013.

[6] H.H. Bauschke and W.M. Moursi. On the Douglas-Rachford algorithm. arXiv:1604.04603, 2016.

[7] R. Beinert and G. Plonka. Ambiguities in one-dimensional phase retrieval of structured functions. *Proc. Appl. Math. Mech.*, 15(1):653–654, 2015.

[8] T. Blumensath and M. Davies. Iterative hard thresholding for compressed sensing. *Appl. Comput. Harmon. A.*, 27(3):265–274, 2009.

[9] T. Blumensath and M. Davies. Normalised iterative hard thresholding; guaranteed stability and performance. *IEEE J. Sel. Topics Signal Process.*, 4(2):298–309, 2010.

[10] Y. M. Bruck and L. G. Sodin. On the ambiguity of the image reconstruction problem. *Opt. Comm.*, 30(3):304–308, Sept. 1979.

[11] J. Buckheit and D. Donoho. Wavelab and reproducible research. In A. Antoniadis et al., editors, *Wavelets and Statistics*. Springer-Verlag, New York, 1995.

[12] E. Candès, Y. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. *SIAM J. Imaging Sci.*, 6(1):199–225, 2013.

[13] E. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inf. Theory*, 51(12):4203–4215, December 2005.

[14] E. Candès and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Inf. Theory*, 52(12):5406–5425, December 2006.

[15] P. L. Combettes and H. J. Trussell. Method of successive projections for finding a common point of sets in metric spaces. *J. Opt. Theory and Appl.*, 67(3):487–507, 1990.

[16] A. L. Dontchev and R. T. Rockafellar. *Implicit Functions and Solution Mapppings*. Springer-Verlag, New York, second edition, 2014.

[17] A. J. J. Drenth, A. Huiser, and H. Ferwerda. The problem of phase retrieval in light and electron microscopy of strong objects. *Opt. Acta*, 22:615–628, 1975.

[18] D. Drusvyatskiy, A. D. Ioffe, and A. S. Lewis. Transversality and alternating projections for nonconvex sets. *Found. Comput. Math.*, 15(6):1637–1651, 2015.

[19] V. Elser. Phase retrieval by iterated projections. *J. Opt. Soc. Amer. A*, 20(1):40–55, 2003.

[20] V. Elser. Strategies for processing diffraction data from randomly oriented particles. *Ultramicroscopy*, 111:788–792, 2011.

[21] J. R. Fienup. Reconstruction of an object from the modulus of its Fourier transform. *Opt. Lett.*, 3(1):27–29, 1978.

[22] D. Gabay. Applications of the method of multipliers to variational inequalities. In *Augmented Lagrangian Methods: Applications to the Solution of Boundary-Value Problems*, pages 299–331. North-Holland, Amsterdam, 1983.

[23] R. W. Gerchberg and W. O. Saxton. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35:237–246, 1972.

[24] M. H. Hayes. The unique reconstruction of multidimensional sequences from Fourier transform magnitude or phase. In H. Stark, editor, *Image Recovery: Theory and Application*. Academic Press, New York, 1987.

[25] M. H. Hayes and A. V. Oppenheim. Signal reconstruction from phase or magnitude. *IEEE Trans. Acoustics, Speech, and Signal Process.*, ASSP-28(6):672–680, December 1980.

[26] R. Hegerl and W. Hoppe. Dynamische theorie der kristallstrukturanalyse durch elektronenbeugung im inhomogenen primrstrahlwellenfeld. *Ber. Bunsenges. Phys. Chem*, 74(11):1148–1154, 1970.

[27] R. Hesse and D. R. Luke. Nonconvex notions of regularity and convergence of fundamental algorithms for feasibility problems. *SIAM J. Optim.*, 23(4):2397–2419, 2013.

[28] R. Hesse, D. R. Luke, and P. Neumann. Alternating projections and Douglas–Rachford for sparse affine feasibility. *IEEE Trans. Signal. Process.*, 62(18):4868–4881, 2014.

[29] R. Hesse, D. R. Luke, S. Sabach, and M. Tam. The proximal heterogeneous block implicit-explicit method and application to blind ptychographic imaging. *SIAM J. Imaging Sci.*, 8(1):426–457, 2015.

[30] K. Huang, Y. C. Eldar, and N. D. Sidiropoulos. Phase retrieval from 1D Fourier measurements: convexity, uniqueness, and algorithms. *IEEE Trans. Signal Process.*, 64(23):6105–6117, 2016.

[31] A. Huiser and H. Ferwerda. The problem of phase retrieval in light and electron microscopy of strong objects, II: on the uniqueness and stability of object reconstruction procedures using two defocused images. *Optica Acta*, 23:445–456, 1976.

[32] K. Jaganathan, Y. C. Eldar, and B. Hassibi. Phase retrieval: an overview of recent developments. arXiv:1510.07713, 2015.

[33] A. Y. Kruger, D. R. Luke, and Nguyen H. T. Set regularities and feasibility problems. *Math. Programming B*, 2016. arXiv:1602.04935.

[34] A. S. Lewis, D. R. Luke, and J. Malick. Local linear convergence of alternating and averaged projections. *Found. Comput. Math.*, 9(4):485–513, 2009.

[35] A. S. Lewis and J. Malick. Alternating projections on manifolds. *Math. Oper. Res.*, 33:216–234, 2008.

[36] S. Loock and G. Plonka. Phase retrieval for Fresnel measurements using a shearlet sparsity constraint. *Inverse Probl.*, 30, 055005, 2014.

[37] S. Loock and G. Plonka. Iterative phase retrieval with sparsity constraints. *PAMM*, 16(1):835–836, 2016.

[38] D. R. Luke. Relaxed averaged alternating reflections for diffraction imaging. *Inverse Probl.*, 21:37–50, 2005.

[39] D. R. Luke. Finding best approximation pairs relative to a convex and a prox-regular set in Hilbert space. *SIAM J. Optim.*, 19(2):714–739, 2008.

[40] D. R. Luke. Local linear convergence of approximate projections onto regularized sets. *Nonlinear Anal.*, 75:1531–1546, 2012.

[41] D. R. Luke. ProxToolbox, 2016. http://num.math.uni-goettingen.de/proxtoolbox/.

[42] D. R. Luke, J. V. Burke, and R. G. Lyon. Optical wavefront reconstruction: theory and numerical methods. *SIAM Rev.*, 44:169–224, 2002.

[43] D. R. Luke, Nguyen H. Thao, and M. T. Tam. Quantitative convergence analysis of iterated expansive, set-valued mappings. http://arxiv.org/abs/1605.05725, May 2016.

[44] A. M. Maiden and J. M. Rodenburg. An improved ptychographical phase retrieval algorithm for diffractiveimaging. *Ultramicroscopy*, 109:1256–1262, 2009.

[45] 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*, 400:342–344, 1999.

[46] B.S. Mordukhovich. *Variational Analysis and Generalized Differentiation, I: Basic Theory; II: Applications*. Grundlehren Math. Wiss. Springer-Verlag, New York, 2006.

[47] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM J. Comput.*, 24:227–234, 1995.

[48] R. Neutze, R. Wouts, D. van der Spoel, E. Weckert, and J. Hajdu. Potential for biomolecular imaging with femtosecond x-ray pulses. *Nature*, 406:752–757, 2000.

[49] D. Noll and A. Rondepierre. On local convergence of the method of alternating projections. *Found. Comput. Math.*, 16(2):425–455, 2016.

[50] A.V. Oppenheim and R.W. Schafer. *Digital Signal Processing*. Prentice Hall, Englewood Cliffs, NJ, 1975.

[51] H. Phan. Linear convergence of the Douglas–Rachford method for two closed sets. *Optimization*, 65:369–385, 2016.

[52] T. Quatieri and A.V. Oppenheim. Iterative techniques for minimum phase signal reconstruction from phase or magnitude. *IEEE Trans. Acoust., Speech, Signal Process.*, ASSP-29(6):1187–1193, Dec. 1981.

[53] J. W. Strutt (Lord Rayleigh). On the interference bands of approximately homogeneous light; in a letter to Prof. A. Michelson. *Phil. Mag.*, 34:407–411, 1892.

[54] V. Stodden, M. McNutt, D. H. Bailey, E. Deelman, Y. Gil, B. Hanson, M. A. Heroux, J. P. A. Ioannidis, and M. Taufer. Enhancing reproducibility for computational methods. *Science*, 354(6317):1240–1241, 2016.

[55] P. Thibault, M. Dierolf, O. Bunk, A. Menzel, and F. Pfeiffer. Probe retrieval in ptychographic coherent diffractive imaging. *Ultramicroscopy*, 109:338–343, 2009.

[56] L. Tian, X. Li, K. Ramchandran, and L. Waller. Multiplexed coded illumination for Fourier ptychography with an LED array microscope. *Biomedical Optics Express*, 5(7):2376–2389, 2014.

[57] P. van Toorn and H. Ferwerda. The problem of phase retrieval in light and electron microscopy of strong objects, III: development of methods for numerical solution. *Opt. Acta*, 23:456–468, 1976.

[58] J. W. Wood, M. A. Fiddy, and R. E. Burge. Phase retrieval using two intensity measurements in the complex plane. *Opt. Let.*, 6(11):514–516, Nov 1981.

[59] G. Zheng, R. Horstmeyer, and C. Yang. Wide-field, high-resolution Fourier ptychographic microscopy. *Nat. Photonics*, 7(9):739–745, 2013.