Algebraic Reconstruction Bounds and Explicit Inversion for Phase Retrieval at the Identifiability Threshold

Franz J. Király
University College London
Department of Statistical Science
Gower Street, London WC1E 6BT

Martin Ehler
University of Vienna
Faculty of Mathematics
Oskar-Morgenstern-Platz 1, A-1090 Vienna

Abstract

We study phase retrieval from rank-one magnitude and more general linear magnitude measurements of an unknown signal as an algebraic estimation problem. It is verified that a certain number of generic rank-one or generic linear measurements are sufficient to enable signal reconstruction for generic signals, and slightly more generic measurements yield reconstructability for all signals. Our results solve few open problems stated in the recent literature. Furthermore, we show how the algebraic estimation problem can be solved by a closed-form algebraic estimation technique, termed ideal regression, providing non-asymptotic success guarantees.

1 INTRODUCTION

Intensity measurements in diffraction imaging, microscopy, and x-ray crystallography represent magnitudes of Fourier samples, and the recovery of their phases is a difficult problem in optical physics. Within a finite model, phase retrieval is the task of reconstructing a vector in $\mathbb{K}^d$ from the magnitude of finitely many rank-1 projections. Classical algorithms are due to Gerchberg/Saxton [14] and Fienup [13] involving alternate projection schemes and fit into standard methods from convex optimization [7], but signal reconstruction is not guaranteed. Sparse nonconvex optimization is applied in [2]. Semidefinite programming is used in [10], but success guarantees are only obtained asymptotically with growing dimension. Algebraic reconstruction formulas were derived in [5], but require the number of measurements to scale quadratically with the dimension. Jointly, algebraic reconstruction and semidefinite programs were applied in [1] to treat rank-k projectors. For further approaches rooted in signal processing, we refer to [12, 20] and references therein. To successfully reconstruct, measurements must contain sufficient information about the signal. If the number of rank-one magnitude measurements is sufficiently large, then generic measurements allow identifiability of all signals, and there is a range of fewer measurements, in which at least generic signals can still be identified, cf. [4]. Measurements using orthogonal projectors of arbitrary rank have been discussed in [9], from where we cite the following open problems:

Appearing in Proceedings of the 17^{th} International Conference on Artificial Intelligence and Statistics (AISTATS) 2014, Reykjavik, Iceland. JMLR: W&CP volume 33.

(1) What is the minimal number of orthogonal projectors enabling phase retrieval for all signals in the real case?
(2) Do sufficiently many generic orthogonal projectors enable phase retrieval for all signals in the real case?
(3) Does the minimal number of required orthogonal projectors for retrieving phases for all signals in the complex case depend on the rank of the projectors?

In view of investigating the above mentioned transition range from generic to identifiability of all signals, we derive three additional questions

(4-6) by replacing “for all signals” in (1-3) with “generic signals”.

The results in [3, 8] directly lead to one more question, which is formulated as a conjecture in [6]:

(7) Do $4n-4$ generic rank-one measurements allow phase retrieval for all signals in the complex case?

So besides the aim for a better understanding of the structure of phase retrieval in general, we are also left with 7 open problems that we intend to solve. In this paper, we claim that phase retrieval is in its core an algebraic problem and emphasize the potential of algebraic tools. This change of perspective enables us to not only answer all of the 7 above questions, but we can also apply symbolic computations and schemes from approximate algebra to design a reconstruction algorithm. Indeed, we observe that phase retrieval can be tackled by ideal regression as introduced in [18] leading to an algebraic signal reconstruction algorithm for few measurements with nonasymptotic success guarantees.

Notes
An extended version of this manuscript is available as [17]. After submission of this paper, question 7 has independently been answered in [11] by different techniques.

Acknowledgments
ME is funded by the Vienna Science and Technology Fund (WWTF) through project VRG12-009. FK is supported by Mathematisches Forschungsinstitut Oberwolfach (MFO).

2 THE ALGEBRA OF PHASE RETRIEVAL

2.1 Algebraization of Phase Retrieval

In this section, we will describe how phase retrieval can be viewed as an algebraic problem. This will be crucial in deriving algebraic solution techniques for phase retrieval. In the usual formulation, the two variants of phase retrieval pose two differently flavoured major obstacles to amenability for algebraic tools: in the real formulation, the mapping is algebraic, but the ground field, the real numbers $\mathbb{R}$, is not algebraically closed. In the complex formulation, the ground field $\mathbb{C}$ is algebraically closed, but the measurement mapping includes complex conjugation, making it non-algebraic. The latter problem can be overcome - as it has been demonstrated for example in [4], by treating the real and imaginary part separately, making the mapping algebraic, but the ground field real in its stead, and therefore reducing the second problem to the first one. We overcome this obstacle by again regarding the algebraic mapping over the complex numbers as base field, and restricting back to the reals when necessary. This procedure allows us to algebraize the measurement process, derive theoretical bounds on reconstructability, and develop accurate reconstruction algorithms. First we recapitulate the measurement process:

Problem 2.1 (Phase Retrieval, original version)**.** Let $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$. Let $z \in \mathbb{K}^n$ be an unknown vector. Let $P_1, \ldots, P_k \in \mathbb{K}^{r\times n}$ be known matrices. Reconstruct z from the measurements

$$b_i = \|P_i z\|^2 = \mathrm{Tr}(zz^* \cdot P_i^* P_i), \quad 1 \le i \le k,$$

and the knowledge of the P_i.

In the usual phase retrieval scenario, the P_i are projectors of rank one. The slightly generalized setting above can be treated with the same mathematical and algorithmical tools, so it means no loss of generality or specificity. Also note that if $\mathbb{K} = \mathbb{R}$, then z can be reconstructed only up to sign, and if $\mathbb{K} = \mathbb{C}$, then only up to phase. We now reformulate the problem, in order to make it amenable to algebraic tools. First we note that phase retrieval is known to be an inverse problem. That is, there is a so-called forward mapping, which takes the (unknown to the observer) signal z, and outputs the (observed) values b_i. The backward problem is then to obtain z from the b_i. Since z can be obtained only up to sign or phase, this is equivalent to obtaining the matrix $Z = zz^*$. Writing all of this explicitly, we obtain as a reformulation of the original Problem 2.1 the following inverse problem:

Problem 2.2. Let $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$. Consider the forward mapping

$$\begin{aligned}\phi : \left(\mathbb{K}^{r\times n}\right)^k \times \mathbb{K}^{n\times n} &\to \left(\mathbb{K}^{r\times n}\right)^k \times \mathbb{K}^k\\ (P_1, \ldots, P_k, Z) &\mapsto \left(P_1, \ldots, P_k, \mathrm{Tr}(Z \cdot P_1^* P_1), \ldots, \mathrm{Tr}(Z \cdot P_k^* P_k)\right).\end{aligned}$$

Reconstruct $\tau := (P_1, \ldots, P_k, Z)$, given $\phi(\tau)$, and assuming that Z is rank one and Hermitian.

Note that we have deliberately included the P_i in the range and the image of ϕ, in order to mathematically model the fact that the projectors P_i are known to the observer; and for technical reasons - equivalent to the latter - which will become apparent further on. Furthermore, assuming that Z is rank one and Hermitian is equivalent to assuming that $Z = zz^*$ for suitable z, since knowing Z is equivalent to know z up to sign/phase. There are two major difficulties in applying algebraic techniques to Problem 2.2. The first is that (A) the base field is not algebraically closed if $\mathbb{K} = \mathbb{R}$, the second being that (B) the mapping ϕ is not algebraic if $\mathbb{K} = \mathbb{C}$, since it includes complex conjugation. The solution approach for problem (A) is relatively straightforward: since the mapping ϕ includes only transposes, it is algebraic, therefore we consider the same mapping over the complex numbers. Also, we replace the matrices $P_i \in \mathbb{R}^{r\times n}$ by matrices $A_i := P_i^\top P_i$ for reason of convenience:

Problem 2.3. Let $z \in \mathbb{C}^n$ be an unknown vector. Consider the forward mapping

$$\begin{aligned}\phi : \left(\mathbb{C}^{n\times n}\right)^k \times \mathbb{C}^{n\times n} &\to \left(\mathbb{C}^{n\times n}\right)^k \times \mathbb{C}^k\\ (A_1, \ldots, A_k, Z) &\mapsto (A_1, \ldots, A_k, \mathrm{Tr}(Z \cdot A_1), \ldots, \mathrm{Tr}(Z \cdot A_k))\end{aligned}$$

Reconstruct $\tau := (A_1, \ldots, A_k, Z)$, given $\phi(\tau)$, and assuming that Z is symmetric rank one, and that the A_i are symmetric of rank r.

There are now several things to note: first, the map ϕ is algebraic, and range and image are now complex. In particular, the measurements can be complex. Note that we want both Z and A_i to be symmetric, not Hermitian, otherwise the problem would not be algebraic. Most importantly, however, Problem 2.3 is a problem which is a-priori different from Problem 2.2, since we have enlarged image and range. When restricting to reals, we obtain the original phase retrieval Problem 2.2, but there is no a-priori reason to believe that the behavior of the complex variant is fundamentally the same as for the original problem. However, as will turn out, Problem 2.3 is much easier amenable to tools from algebraic geometry, both on the theoretical and the practical side. Results and algorithms will give rise to solutions for questions and tasks over the reals, as it will be explained in the following section. We proceed treating the variant of the phase retrieval problem 2.2 where complex signals are allowed. Recall that the problem was that (B) the map ϕ is not algebraic. The solution for this is to "algebraize" the map by considering real and imaginary part separately. Namely, writing $P_i = Q_i + \iota \cdot S_i$ with $Q_i, S_i \in \mathbb{R}^{m\times n}$ and $z = x + \iota y$, where ι denotes the imaginary unit, we obtain:

Problem 2.4. Let $x, y \in \mathbb{R}^n$ be unknown vectors, write $R := xx^\top + yy^\top$ and $\Phi := yx^\top - xy^\top$. Also, write $B_i := Q_i^\top Q_i + S_i^\top S_i$ and $C_i := Q_i^\top S_i - S_i^\top Q_i$ for $Q_i, S_i \in \mathbb{R}^{m\times n}$. Consider the forward mapping $\phi : (\mathbb{R}^{n\times n})^{2k} \times \mathbb{R}^n \to (\mathbb{R}^{n\times n})^{2k} \times \mathbb{R}^k$, $(B_1, C_1, \dots, B_k, C_k, R, \Phi) \mapsto$

$$(B_1, C_1, \dots, B_k, C_k, \mathrm{Tr}(R\cdot B_1 + \Phi\cdot C_1), \dots, \mathrm{Tr}(R\cdot B_n + \Phi\cdot C_n)).$$

Reconstruct $\tau = (B_1, C_1, \dots, B_k, C_k, R, \Phi)$, given $\phi(\tau)$, assuming that B_i, C_i, R, Φ were of the above form.

An elementary computation shows that Problem 2.4 is equivalent to the original complex phase retrieval problem 2.1: namely, $zz^* = R + \iota\Phi$, so knowing R and Φ is equivalent to knowing z up to phase. Observe that ϕ is now an algebraic map, since the rule is algebraic, and so is the possible set of B_i, C_i, X, Y. However, the mapping ϕ is now over the reals, a field which is not algebraically closed, entailing an analogue of complication (A) which we have treated in the real case by allowing complex matrices in the range. We will once more do the same and allow a complex range. The set of matrices though have a very specific structure, so we introduce notation for them in our final formulation of the complex phase retrieval problem:

Problem 2.5 (algebraized phase retrieval of complex signal)**.** Define the following sets of matrices:

$$\mathcal{S}_\mathbb{C} := \{(xx^\top + yy^\top, yx^\top - xy^\top) : x, y \in \mathbb{C}^n\},$$
$$\mathcal{P}_\mathbb{C}(r) := \{(Q^\top Q + S^\top S, Q^\top S - S^\top Q) : S, Q \in \mathbb{C}^{r\times n}\}.$$

Consider the forward mapping

$$\phi : \mathcal{P}_\mathbb{C}(r)^k \times \mathcal{S}_\mathbb{C} \to \mathcal{P}_\mathbb{C}(r)^k \times \mathbb{C}^k, \quad (B_1, C_1, \dots, B_k, C_k, R, \Phi)$$
$$\mapsto (B_1, C_1, \dots, B_k, C_k, \mathrm{Tr}(R\cdot B_1 + \Phi\cdot C_1), \dots, \mathrm{Tr}(R\cdot B_n + \Phi\cdot C_n)).$$

Given $\tau = \phi(B_1, C_1, \dots, B_k, C_k, R, \Phi)$, determine $\phi^{-1}(\tau)$.

The set $\mathcal{S}_\mathbb{C}$ parameterizes the possible signals, while $\mathcal{P}_\mathbb{C}(r)$ parameterizes the possible projections (of rank r). Note that $\mathcal{S}_\mathbb{C} = \mathcal{P}_\mathbb{C}(1)$; nevertheless we make this notational distinction between $\mathcal{S}_\mathbb{C}$ and $\mathcal{P}_\mathbb{C}(.)$ for clarity. We reformulate the phase retrieval problem for real signals in analogy, by defining symbols for the space of matrices, yielding in the final version:

Problem 2.6 (algebraized phase retrieval of real signal)**.** Define the following sets of matrices:

$$\mathcal{S}_\rho := \{zz^\top : z \in \mathbb{C}^n\}, \quad \mathcal{P}_\rho(r) := \{P_i^\top P_i : P_i \in \mathbb{C}^{r\times n}\}.$$

Consider the forward mapping

$$\phi : \mathcal{P}_R(r)^k \times \mathcal{S}_R \to \mathcal{P}_R(r)^k \times \mathbb{C}^k$$
$$(A_1, \dots, A_k, Z) \mapsto (A_1, \dots, A_k, \mathrm{Tr}(Z\cdot A_1), \dots, \mathrm{Tr}(Z\cdot A_n))$$

Given $\tau = \phi(A_1, \dots, A_k, Z)$, determine $\phi^{-1}(\tau)$.

Observe that $\mathcal{S}_\rho$ models the possible signals, and is exactly the set of symmetric complex matrices of rank 1 (or less), whereas $\mathcal{P}_\rho(r)$ models the projections, and is exactly the set of symmetric complex matrices of rank r (or less). Note that we have formulated both the real and the complex problem with almost the same forward mapping, the difference lies in the different sets of projection matrices, where in the real case we have single matrices, the complex case yields related pairs. Also, for the complex variant of phase retrieval, we have related pairs of matrices R and Φ instead of the single matrix Z. To make the notation uniform for both the real and complex cases, we introduce the following convention:

Notation 2.7. Let $Z, A \in \mathbb{C}^{n\times n} \times \mathbb{C}^{n\times n}$, with $Z = (X, Y)$ and $A = (B, C)$. Then, we will write, by convention, $\mathrm{Tr}(Z\cdot A) := \mathrm{Tr}(X\cdot B + Y\cdot C)$.

2.2 Identifiability and Genericity

A signal z is called *identifiable* if it is uniquely determined in $\mathbb{K}^n$ by the measurements b_i up to a global phase factor, which is an ambiguity one cannot avoid. The choice of k generic measurements by means of rank-1 projectors yield identifiability of generic signals if and only if $k \geq n+1$ in the real and $k \geq 2n$ in the complex case, cf. [4, Theorems 2.9 and 3.4]. Generic rank-1 projectors yield identifiability for all signals if and only if $k \geq 2n-1$ in the real case. For the complex setting, examples with $k \geq 4n-4$ are known, and this bound is conjectured to be necessary [8].

We will generalize the statements to the scenario of general linear projections. As described earlier, the strategy is to consider first the corresponding algebraized problem over an algebraically closed field, namely $\mathbb{C}$, instead of $\mathbb{R}$, and then descend the results back to the real numbers $\mathbb{R}$. Again, it is important to note that this is subtly different from considering the projection problem over the complex numbers, since instead of complex conjugation, we consider transposition in order to keep the problem algebraic.

A Short Note on Technical Conditions

The following exposition will use some technical conditions on varieties and maps, namely them being *irreducible*, and (generically) *unramified*. These are standard notions in algebraic geometry and can be found in most introductory books - we refrain from explaining them here as this is beyond the scope of the paper; the logic in the proofs can be understood without knowing what these mean exactly - a glossary of definitions can be found in Appendix A.1. Intuitively, an algebraic set being irreducible means that there is only one prototypical behaviour for its elements. Unramifiedness is a point-wise algebraic certificate for a mapping staying stable under perturbation in a certain sense. In our case, unramifiedness will certify for identifiability which is stable under perturbation of signals or measurements.

2.2.1 Identifiability of Signals

In this paragraph, we translate identifiability of a signal into an algebraic statement. The main concepts will be identifiability, and identifiability which is stable under perturbation, both corresponding to certain algebraic properties of the signal.

Notation 2.8. We fix some notation and technical assumptions that will be valid in the relevant cases of real and complex phase recognition:

(i) The signals will be modelled by an irreducible variety $\mathcal{S} \subseteq (\mathbb{C}^{n\times n})^{\gamma}$, with $\gamma = 1$ in the real and $\gamma = 2$ in the complex case. For example, $\mathcal{S} = \mathcal{S}_{\rho}$ or $\mathcal{S} = \mathcal{S}_{\mathbb{C}}$, as in Section 2.1.

(ii) A measurement scheme will be modelled by the tuple $A = (A_1, \ldots, A_k) \in \left((\mathbb{C}^{n\times n})^{\gamma}\right)^k$ with $k \in \mathbb{N}$ being the number of measurements.

(iii) The measurement process is the formal mapping $\phi_A : \mathcal{S} \to \mathbb{C}^k, \quad Z \mapsto (\mathrm{Tr}(Z \cdot A_1), \ldots, \mathrm{Tr}(Z \cdot A_k))$.

The condition that $\mathcal{S}$ is irreducible is fulfilled in the cases discussed in the introductory Section 2.1. Namely, both $\mathcal{S}_{\rho}$ and $\mathcal{S}_{\mathbb{C}}$ are irreducible varieties, as it is proved in Proposition B.3. The following statement is crucial in obtaining our local-to-global principle for identifiability. It characterizes signals which are identifiable and stably so under perturbation just in terms of the signal itself, therefore allowing to remove any reference to open neighbourhoods.

Proposition 2.9. *Assume that ϕ_A is generically unramified. Let $Z \in \mathcal{S}$. Then, the following three statements are equivalent:*

(i) *Z is identifiable from $\phi_A(Z)$, and remains identifiable under infinitesimal perturbation*[1].

(ii) *Z is identifiable from $\phi_A(Z)$, and ϕ_A is unramified over Z.*

(iii) *A generic*[2] *$Y \in \mathcal{S}$ is identifiable from $\phi_A(Z)$.*

Intuitively, Proposition 2.9 means that an identifiable signal which remains so under perturbation certifies for the whole signal space. It is also important to note that condition (ii) in Proposition 2.9 is essentially independent from the choice of $\mathcal{S}$ while (i) and (iii) are a-priori not. We introduce terminology for the condition described in (i):

Definition 2.10. For brevity, we will call a signal $Z \in \mathcal{S}$ that is identifiable from $\phi_A(Z)$, and remains identifiable under infinitesimal perturbation, a *perturbation-stably identifiable* signal.

We can reformulate Proposition 2.9 as a principle of excluded middle, stating that either almost all signals are perturbation-stably identifiable, or none:

Corollary 2.11. **(i)** *If there exists a signal $Z \in \mathcal{S}$ which is perturbation-stably identifiable from $\phi_A(Z)$, then a random signal $Y \in \mathcal{S}$ is perturbation-stably identifiable with probability one under any Hausdorff continuous probability density on $\mathcal{S}$.*

(ii) *It cannot happen that there are sets $\mathcal{A}, \mathcal{B} \subseteq \mathcal{S}$, both with positive Hausdorff measure, such that all signals $Z \in \mathcal{A}$ are perturbation-stably identifiable, and all signals $Z \in \mathcal{A}$ are not perturbation-stably identifiable.*

2.2.2 Identifyingness as a Measurement Property

In Corollary 2.11, it has been shown that if one signal is perturbation-stably identifiable, then almost all signals are. Therefore the fact whether almost all signals are identifiable can be regarded as a property of the measurement regime. The following theorem makes this statement exact and states that measurement regimes fall into exactly one of three classes:

Theorem 1. *For a fixed measurement regime $(A_1, \ldots, A_k)$, consider the three cases*

(a) *A generic signal $Z \in \mathcal{S}$ is not identifiable from $\phi_A(Z)$.*

(b) *A generic, but not all signals $Z \in \mathcal{S}$, are identifiable from $\phi_A(Z)$.*

(c) *All signals $Z \in \mathcal{S}$ are identifiable from $\phi_A(Z)$.*

The three cases above are mutually exclusive and exhaustive, and equivalent to

[1] That is, there is a relatively Borel-open neighborhood $U \subseteq \mathcal{S}$ with $Z \in U$ such that for all $Y \in U$, it holds that $\#\phi_A^{-1}\phi_A(Z) = 1$.

[2] That is, the set of non-identifiable $Y \in \mathcal{S}$ is a proper Zariski closed subset and therefore Hausdorff measure zero subset of $\mathcal{S}$.

(a) *No signal $Z \in \mathcal{S}$ is perturbation-stably identifiable from $\phi_A(Z)$.*
(b) *A generic, but not all signals $Z \in \mathcal{S}$, are perturbation-stably identifiable from $\phi_A(Z)$.*
(c) *All signals $Z \in \mathcal{S}$ are perturbation-stably identifiable from $\phi_A(Z)$.*

A proof is provided in Appendix B.3. Theorem 1 allows to regard the different grades of identifiability (a), (b), (c) as properties of the measurement regime. We therefore introduce the following abbreviating notation:

Definition 2.12. We call a measurement tuple $A = (A_1, \dots, A_k)$:

(a) *non-identifying* for signals in $\mathcal{S}$, if no signal $Z \in \mathcal{S}$ is perturbation-stably identifiable from $\phi_A(Z)$.
(b) *generically identifying* for signals in $\mathcal{S}$, if generic signals $Z \in \mathcal{S}$ are (perturbation-stably) identifiable from $\phi_A(Z)$, and *incompletely identifying*, if generic, but not all signals $Z \in \mathcal{S}$ are (perturbation-stably) identifiable from $\phi_A(Z)$.
(c) *completely identifying* for signals in $\mathcal{S}$, if all signals $Z \in \mathcal{S}$ are (perturbation-stably) identifiable from $\phi_A(Z)$.

Theorem 1 then can be rephrased that a measurement regime $A_1, \dots, A_k$ is either non-identifying, incompletely identifying, or completely identifying - note that due to the theorem, it does not matter whether the "perturbation-stably" in the brackets is there or not. We now show that these are properties of the space of possible measurements, just as identifiability is not only a property of the signal, but of signal space.

Notation 2.13. We introduce some notation modelling the space of measurements:

(iv) The space of measurements of type $(A_1, \dots, A_k)$ will be modelled by irreducible varieties $\mathcal{P}_1, \dots, \mathcal{P}_k \subseteq (\mathbb{C}^{n \times n})^\gamma$, with $\gamma = 1$ in the real and $\gamma = 2$ in the complex case. We will write $\mathcal{P}^{(k)} = \mathcal{P}_1 \times \dots \times \mathcal{P}_k$ for the space of measurement tuples of size k. For example, $\mathcal{P}^{(k)} = \mathcal{P}_\mathbb{C}(r)^k$ for complex signals, or $\mathcal{P}^{(k)} = \mathcal{P}_\rho(r)^k$ for real ones.
(v) The extended measurement process will be modelled by the formal forward mapping

$$\begin{aligned} \phi : \mathcal{P}^k \times \mathcal{S} &\to \mathcal{P}^k \times \mathbb{C}^k \\ (A_1, \dots, A_k, Z) &\mapsto (A_1, \dots, A_k, \operatorname{Tr}(Z \cdot A_1), \dots, \operatorname{Tr}(Z \cdot A_n)). \end{aligned}$$

The condition that the $\mathcal{P}_i$ is irreducible is fulfilled in the cases discussed in the introductory Section 2.1: both $\mathcal{P}_\rho(r)$ and $\mathcal{P}_\mathbb{C}(r)$ are irreducible varieties, see Proposition B.3. Our main result is an analogue to the characterization in Proposition 2.9, now for the measurement matrices:

Proposition 2.14. *Assume that ϕ is generically unramified. Then, the following three statements are equivalent:*

(i) *$Z \in \mathcal{S}$ is identifiable from $\phi(A, Z)$, and remains identifiable under infinitesimal perturbation*[3] *of A and Z.*
(ii) *(A, Z) is identifiable from $\phi(A, Z)$, and ϕ is unramified over (A, Z).*
(iii) *For generic $B \in \mathcal{P}^k$, a generic $Y \in \mathcal{S}$ is identifiable*[4] *from $\phi(A, Y)$.*

In particular, condition (i) is a Zariski open property on the measurement-signal-pair (A, Z); that is, the set of measurement-signal-pairs (A, Z) with property (i) is a Zariski open subset of $\mathcal{P}^k \times \mathcal{S}$.

The main obstacle in generalizing Theorem 1 to an algebraic characterization, or a local-global-property of measurements lies in the fact that the perturbation can occur in both the signal Z and the measurement regime A. We therefore need to provide an intermediate result which removes the dependence on the measurement:

Proposition 2.15. *Assume that ϕ is generically unramified. Then, the following two conditions on measurement regimes $A \in \mathcal{P}^k$ are (Zariski) open conditions:*

(i) *A is generically identifying and remains generically identifying under perturbation. That is, there is a (relatively Borel-) open neighborhood $U \subseteq \mathcal{P}^k$ with $A \in U$ such that all $B \in U$ are generically identifying.*
(ii) *A is completely identifying and remains completely identifying under perturbation. That is, there is a (relatively Borel-) open neighborhood $U \subseteq \mathcal{P}^k$ with $A \in U$ such that all $B \in U$ are completely identifying.*

A proof is given in Appendix B.3.

Definition 2.16. We call a measurement regime $A \in \mathcal{P}^{(k)}$:

(a) *stably non-identifying* in $\mathcal{P}^{(k)}$, if A is non-identifying and remains non-identifying under perturbation, as in Proposition 2.15 (i).
(b) *stably generically identifying* in $\mathcal{P}^{(k)}$, if A is generically identifying and remains generically identifying under perturbation, as in condition (i). *stably incompletely identifying* in $\mathcal{P}^{(k)}$, if A is incompletely identifying and remains incompletely identifying under perturbation, as in Proposition 2.15 (i).
(c) *stably completely identifying* in $\mathcal{P}^{(k)}$, if A is completely identifying and remains completely identifying under perturbation, as in Proposition 2.15 (ii).

[3] That is, there is a relatively Borel-open neighborhood $U \subseteq \mathcal{P}^k \times \mathcal{S}$ with $(A, Z) \in U$ such that for all $Y \in U$, it holds that $\#\phi^{-1}\phi(Y) = 1$.

[4] That is, the set of $(B, Y) \in \mathcal{P}^k \times \mathcal{S}$ where $Y \in \mathcal{S}$ is non-identifiable from $\phi(B, Y)$ is a proper Zariski closed subset and therefore Hausdorff measure zero subset of $\mathcal{P}^k \times \mathcal{S}$.

If $\mathcal{P}^{(k)}$ is obvious from the context, we will omit the qualifier "in $\mathcal{P}^{(k)}$", always keeping in mind that the terminology depends on $\mathcal{P}^{(k)}$.

Proposition 2.15 allows to prove an analogue of Theorem 1, now for classes of measurements instead of a single measurement regime:

Theorem 2. *Assume that ϕ is generically unramified. Consider the three cases*

- **(a)** *A generic measurement regime $A \in \mathcal{P}^k$ is non-identifying.*
- **(b)** *A generic measurement regime $A \in \mathcal{P}^k$ is incompletely identifying.*
- **(c)** *A generic measurement regime $A \in \mathcal{P}^k$ is completely identifying.*

The three cases above are mutually exclusive and exhaustive, and equivalent to

- **(a)** *A generic measurement regime $A \in \mathcal{P}^k$ is stably non-identifying. No measurement regime $A \in \mathcal{P}^k$ is stably generically identifying.*
- **(b)** *A generic measurement regime $A \in \mathcal{P}^k$ is stably incompletely identifying.*
- **(c)** *A generic measurement regime $A \in \mathcal{P}^k$ is stably completely identifying.*

We can therefore define terminology that describe cases (a) to (c) shortly:

Definition 2.17. Keep the notations of Theorem 2. We will call a the set of measurements $\mathcal{P}^k$ *generically unramified* if ϕ is generically unramified. We will call a generically unramified $\mathcal{P}^k$:

- **(a)** *non-identifying* if a generic measurement $A \in \mathcal{P}^k$ is non-identifying.
- **(b)** *generically identifying* if a generic measurement $A \in \mathcal{P}^k$ is generically identifying. *incompletely identifying* if a generic measurement $A \in \mathcal{P}^k$ is incompletely identifying.
- **(c)** *completely identifying* if a generic measurement $A \in \mathcal{P}^k$ is completely identifying.

2.3 Transfer Results for Identifyingness

In this section we will collect different results that allow to transfer identifiyingness properties from one set of potential measurements to another. Proofs can be found in Appendix B.3.

Notation 2.18. We will consider irreducible varieties $\mathcal{P}^{(k)} = \mathcal{P}_1 \times \cdots \times \mathcal{P}_k$ and $\mathcal{Q}^{(k)} = \mathcal{Q}_1 \times \cdots \times \mathcal{Q}_k$, with corresponding forward maps ϕ, φ.

Lemma 2.19. *Assume $\mathcal{P}_i \subseteq \mathcal{Q}_i$ for all i, that is, $\mathcal{P}^{(k)} \subseteq \mathcal{Q}^{(k)}$. Then: if $\mathcal{P}^{(k)}$ is generically unramified/generically identifying/completely identifying, then so is $\mathcal{Q}^{(k)}$.*

Lemma 2.20. *Assume the $\mathcal{P}_i \subseteq (\mathbb{C}^{n\times n})^\gamma$ are all spaces of rank at most r_i matrices, that is, of the form $\mathcal{P}_\rho(r_i)$ or $\mathcal{P}_\mathbb{C}(r_i)$. Assume that the $\mathcal{Q}_i$ are the corresponding variety of orthogonal/unitary projection matrices of rank exactly r_i. Then, $\mathcal{P}^{(k)}$ is generically/completely identifying if and only if $\mathcal{Q}^{(k)}$ is.*

In our terminology, Proposition 2.15 also implies that the behavior of random projectors is completely determined by their number, and no other properties. This motivates the following:

Definition 2.21. Consider a family of irreducible varieties $\mathcal{P} = \{\mathcal{P}_i\}_{i\in\mathbb{N}}$. We will denote the smallest number k such that

- **(i)** $(\mathcal{P}_1, \dots, \mathcal{P}_k)$ is generically identifying by $\lambda(\mathcal{P}_1, \mathcal{P}_2, \dots) = \lambda(\mathcal{P})$ and call it the *generic identifiability threshold*.
- **(ii)** $(\mathcal{P}_1, \dots, \mathcal{P}_k)$ is completely identifying by $\kappa(\mathcal{P}_1, \mathcal{P}_2, \dots) = \kappa(\mathcal{P})$ and call it the *complete identifiability threshold*.

If $\mathcal{P}_i = \mathcal{X}$ for all i, for some variety $\mathcal{X}$, we also write $\lambda(\mathcal{X})$ and $\kappa(\mathcal{X})$ instead of $\lambda(\mathcal{P}_1, \mathcal{P}_2, \dots)$ and $\kappa(\mathcal{P}_1, \mathcal{P}_2, \dots)$.

2.4 Identifiability of Real Signals

We derive bound on the identifiability thresholds for real signals. Proofs can be found in Appendix B.3.

Proposition 2.22. *Consider identifiability from real signals $\mathcal{S} = \{zz^\top, z \in \mathbb{C}^n\}$. For any family of irreducible varieties $\mathcal{P}_i \subseteq \mathbb{C}^{n\times n}, i \in \mathbb{N}$, with $n \geq 2$, it holds that $\kappa(\mathcal{P}) \geq \lambda(\mathcal{P})$, and $\lambda(\mathcal{P}) \geq n+1$.*

Some bounds for real signals can be readily inferred from literature:

Theorem 3. *Consider identifiability from real signals, corresponding to the complex signal variety $\mathcal{S}_\rho = \{zz^\top, z \in \mathbb{C}^n\}$, and projectors $\mathcal{P} = \mathcal{S}$. Then:*
$\lambda(\mathcal{P}) = n+1, \quad$ *and* $\quad \kappa(\mathcal{P}) = 2n-1.$

By virtue of Lemma 2.19, these results can immediately be broadened to include general linear projections, while Lemma 2.20 yields the case of orthogonal measurements:

Theorem 4. *Consider identifiability from real signals, corresponding to the complex signal variety $\mathcal{S} = \{zz^\top, z \in \mathbb{C}^n\}$, and the family $\mathcal{P}_i = \{P^\top \cdot P \ : \ P \in \mathbb{C}^{r_i\times n}\}, i \in \mathbb{N}$ of projectors of potentially different ranks $r_i \geq 1$. Then:*
$\lambda(\mathcal{P}) = n+1, \quad$ *and* $\quad \kappa(\mathcal{P}) = 2n-1.$ *The result remains unaltered if the projectors $\mathcal{P}$ are restricted to be orthogonal.*

Using the tools introduced in Section B.2, we obtain from this statement about the complexified problem one about the original phase retrieval problem for the reals:

Theorem 5. *Let $P_i \in \mathbb{R}^{r_i \times n}, 1 \leq i \leq k$ be generic. Then, a generic signal $z \in \mathbb{R}^n$ is identifiable from $b_i = \|P_i z\|^2, 1 \leq i \leq k$ up to sign if and only if $k \geq n+1$. All signals $z \in \mathbb{R}^n$ are identifiable from $b_i = \|P_i z\|^2, 1 \leq i \leq k$ up to sign if and only if $k \geq 2n-1$. The result remains unaltered if the projectors P_i are restricted to be orthogonal.*

This solves the open problems (1-6).

2.5 Identifiability of Complex Signals

The case of complex phase recognition is somewhat analogous to the real one, while more technical due to the special structure of the matrices involved. The proof logic is analogous to the case of real signals, we combine results from literature with our own bounds and transfer statements to obtain theorems. The complete proofs and theorems are in Appendix B.4.

Theorem 6. *Let $P_i \in \mathbb{C}^{r_i \times n}, 1 \leq i \leq k$ be generic. Then, a generic signal $z \in \mathbb{C}^n$ is identifiable from $b_i = \|P_i z\|^2, 1 \leq i \leq k$ up to phase if and only if $k \geq 2n$. All signals $z \in \mathbb{C}^n$ are identifiable from $b_i = \|P_i z\|^2, 1 \leq i \leq k$ up to phase if $k \geq 4n-4$. The result remains unaltered if the projectors P_i are restricted to be unitary.*

This solves problem (7), and problems (1-6) for unitary projection matrices.

3 ALGEBRAIC INVERSION

3.1 Phase Retrieval as Ideal Regression

We will show that the phase retrieval problem is a special case of an algebraic estimation problem, called ideal regression. This means that not only is the solvability and identifiability of the problem determined by algebraic invariants, such as n, k, or the kind of projectors, but that it is - in principle - also accessible to algorithmical estimation tools from approximate algebra, such as those presented in [18], yielding explicit and deterministic inversion formulae not only for $k = \Omega(n^2)$, but directly at the identifiability threshold $k \geq n+1$. The reformulation of the phase retrieval as an algebraic estimation problem bears similarities to the algebraization in Section 2.1. The major idea consist of converting the observation into polynomials, which are then manipulated to obtain the solution. Assume we are in the case of the real phase recognition problem, wanting to identify a signal $z \in \mathbb{R}^n$. Then, let $X = (X_1, \dots, X_n)$ be a vector of formal variables. The k projection matrices P_i give rise to k polynomials

$$p_i(X_1, \dots, X_n) = X^\top A_i X - b_i$$

in the variables X_j, with $A_i = P_i^\top P_i$, such that, after substitution, we have $p_i(z) = 0$. By definition the polynomials p_i are contained in the ideal $\mathfrak{I} := \mathrm{I}(z) \subseteq \mathbb{C}[X_1, \dots, X_n]$. Thus, the estimation problem becomes, for the real phase recognition problem:

Problem 3.1. Let $z \in \mathbb{R}^n$ be unknown, let $\mathfrak{s} = \langle X_1 - z_1, \dots, X_n - z_n \rangle \in \mathbb{C}[X_1, \dots, X_n]$. Let $p_1, \dots, p_k \in \mathfrak{I}$ be known polynomials, of the form $p_i(X_1, \dots, X_n) = X^\top A_i X - b_i$, where $b_i = z^\top A_i z - b_i$. Then, reconstruct $\mathfrak{s}$, or equivalently, z, from the polynomials $p_1, \dots, p_k, 1 \leq i \leq k$.

What at first seems like a mere reformulation, contains the gist of the algebraic ideal regression method: instead of fitting a loss function or performing optimization on z, or taking the b_i, P_i as an input, we try to obtain the solution from manipulating the polynomials p_i as symbolic objects in their own right. Again, we note that we are working over the complex numbers in the polynomial ring $\mathbb{C}[X_1, \dots, X_n]$, similarly to the algebraization; we will again show that this is no major problem, from an algorithmic aspect. The complex case is slightly different but can be treated similarly. Here, let $X = (X_1, \dots, X_n)$ and $Y = (Y_1, \dots, Y_n)$ be vectors of formal variables, and let $P_i = Q_i + \iota \cdot S_i$ with $Q_i, S_i \in \mathbb{R}^{m \times n}$. The projections give rise to k polynomials

$$p_i = (X, Y)^\top \begin{pmatrix} Q_i^\top Q_i + S_i^\top S_i & S_i^\top Q_i - Q_i^\top S_i \\ Q_i^\top S_i - S_i^\top Q_i & Q_i^\top Q_i + S_i^\top S_i \end{pmatrix} \begin{pmatrix} X \\ Y \end{pmatrix},$$

and those are, similar to the real case, contained in the ideal $\mathfrak{I} := \mathfrak{s}((X, Y) - \tilde{z}) \subseteq \mathbb{C}[X_1, \dots, X_n, Y_1, \dots, Y_n]$, where $\tilde{z} = (\Re z, \Im z) \in \mathbb{R}^{2n}$. So the estimation problem is, in the complex case:

Problem 3.2. Let $\tilde{z} \in \mathbb{R}^{2n}$ be an unknown point, let $\mathfrak{s} = \langle X_1 - \Re z_1, Y_1 - \Im z_1, \dots, X_n - \Re z_n, Y_n - \Im z_n \rangle \subseteq \mathbb{C}[X_1, \dots, X_n, Y_1, \dots, Y_n]$. Let $p_1, \dots, p_k \in \mathfrak{s}$ be known polynomials, of the form as above. Reconstruct $\mathfrak{s}$, or equivalently $\tilde{z}$, from the $p_1, \dots, p_k, 1 \leq i \leq k$.

Note that the ideal regression formulation of phase retrieval Problem 3.2 differs fundamentally from the algebraized inverse problem version given in Problem 2.5, since in ideal regression, we split real and complex parts of the formal variables, whereas in the algebraization, we split real and complex parts of the matrices involved. Still, both problems are intrinsically related, and can be considered, in a certain sense, as each other's duals.

3.2 An Inversion Formula with Ideal Regression

We describe how the ideal regression formulation of the phase retrieval problem 3.1 can be solved by an approximate algebraic algorithm; we focus on the real case. If $k \geq \binom{n+1}{2}$, there exist explicit inversion formulae in which

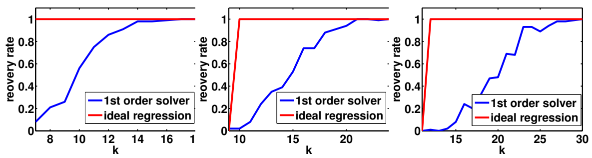

(a) $n = 6$ (b) $n = 8$ (c) $n = 10$

Figure 1: Recovery rates averaged over 100 repeats without any noise for ideal regression and for a first order solver in PhaseLift.

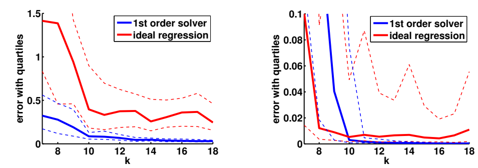

(a) Noise level $\sigma = 10^{-2}$. (b) Noise level $\sigma = 10^{-4}$.

Figure 2: Mean squared error for $n = 6$ and quartiles for 100 repeats.

one computes an approximation for $\mathrm{Tr}(A_i zz^\top) = b_i$, which is now considered as a linear system of k equations in the $\binom{n+1}{2}$ unknowns $zz^\top$; this can be written as pseudo-inverting a matrix which has one row per A_i, and noise stability can be achieved by regularization. If $k \quad \binom{n+1}{2}$, such a direct approach will not work. However, it is nevertheless possible to construct an explicit deterministic inversion formula, readily providing answers at the identifiability threshold $k \geq n+1$, and which is numerically stable. The main idea is to use an ideal regression algorithm, namely Algorithms 1 and 2 in [18]; the ideal $\mathfrak{s}$ we wish to estimate in our case is linear, namely $\mathfrak{s} = \langle X_1 - z_1, \ldots, X_n - z_n \rangle$, and the input polynomials are of degree two, contained in $\mathfrak{s}$. Since $\mathfrak{s}$ is inhomogenous, Algorithm 1 in [18] will output the homogenous part of $\mathfrak{s}$, namely $\mathfrak{s}_h = \mathfrak{s} \cap \langle X_1, \ldots, X_n \rangle$ which is also linear, and can be used to estimate z. Instantiating Algorithm 1 in [18] with $D = n, d = 1$, and polynomials $f_i := p_i/b_i - \overline{p}, 1 \leq i \leq k-1$, where $\overline{p} = \sum_{i=1}^k p_i/b_i$, yields an estimate for generators ℓ_1, ℓ_{n-1} of $\mathfrak{s}_h$. The signal z fulfills $\ell_i(z) = 0$, therefore z is orthogonal to the coefficient vectors of the ℓ_i and can be determined up to a scalar multiple $z' = \alpha z$ from the ℓ_i. Thus, z can be determined by setting $z := z'/\alpha$ where α can be estimated as $\alpha := \exp\left(\sum_{i=1}^k \log\left((z')^\top A_i z\right) - \log b_i\right)$. We will refer to this strategy as the "explicit inversion" in the experiments section. We refrain from actually explaining in detail how Algorithm 1 in [18] works, or from stating the algorithm itself, due to the amount of notational overhead which would be needed, and refer the reader to the original paper instead. We want to stress that Algorithm 1 is deterministic and numerically stable, therefore it yields a potentially explicit and regularizable inversion formula for the phase recognition problem.

4 EXPERIMENTS

In this section we provide few numerical experiments illustrating that generic real signals can be identified from few generic magnitude measurements by using the inversion formula obtained from ideal regression as outlined in section 3.2. We also include a few comparisons to an alternative method. Classical phase retrieval algorithms such as Gerchberg/Saxton [14] and Fienup's alternatives [13] are customized to Fourier measurements, hence are also limited to this setting. An approach that can deal with generic measurements is PhaseLift [10], which is based on finding the feasible point of a semidefinite program and is proposed to be solved using first order methods. The theoretical results in [10] are asymptotic in the ambient dimension n and no success guarantees are derived for fixed n. Nonetheless, PhaseLift is known to be quite successful and very robust against noise in practise. The complexity of ideal regression causes limits in the number of measurements that can be dealt with in practise, while it yields an explicit reconstruction formula. We shall study the performance of ideal regression and PhaseLift for few measurements. In the numerical experiments, we choose the signal x uniformly distributed on the sphere. Measurements are performed by orthogonal rank-1 projectors, also uniformly distributed (according to the standard Haar measure on this set), and we deal with corrupted measurements $\tilde{b} = b + \eta$, where η is Gaussian white noise of variance σ. The outcome of performance comparisons between ideal regression and PhaseLift very much depend on the noise level. If measurements are exact, then ideal regression yields signal recovery for generic $n+1 \leq k \leq 3n$ measurements, a range, in which PhaseLift performs rather poorly, see Fig. 1 for $n = 6, 8, 10$. For inexact yet still very accurate measurements, in other words very low noise levels ($\sigma \approx 10^{-4}$), ideal regression still outperforms PhaseLift when the number of measurements is close to the threshold $n+1$, see Fig. 2(b), with a comparable accuracy for higher noise levels ($\sigma \approx 10^{-2}$), cf. Fig. 2(a). Nonetheless, it must be mentioned that with slightly larger and hence more common noise levels, especially when the number of measurements increases, then PhaseLift is eventually to be favored since error rates are then significantly smaller than within ideal regression. It is interesting to note that ideal regression performs well close to the identifiability threshold $k = n+1$, whereas PhaseLift yields more accurate estimates as the number of samples increases.

References

[1] Christine Bachoc and Martin Ehler. Signal reconstruction from the magnitude of subspace components. *arXiv e-prints*, September 2012. arXiv:1209.5986.

[2] R. Balan, P. Casazza, and D. Edidin. Equivalence of reconstruction from the absolute value of the frame coefficients to a sparse representation problem. *IEEE Signal Process. Lett.*, 14(5):341–343, 2007.

[3] Radu Balan. Stability of phase retrievable frames. *arXiv e-prints*, August 2013. arXiv:1308.5465.

[4] Radu Balan, Pete Casazza, and Dan Edidin. On signal reconstruction without phase. *Appl. Comput. Harmon. Anal*, 20:345–356, 2006.

[5] Radu Balan, Bernhard G. Bodmann, Peter G. Casazza, and Dan Edidin. Painless reconstruction from magnitudes of frame coefficients. *J. Fourier Anal. Appl.*, 15(4):488–501, 2009.

[6] Afonso S. Bandeira, Jameson Cahill, Dustin G. Mixon, and Aaron A. Nelson. Saving phase: Injectivity and stability for phase retrieval. *arXiv e-prints*, February 2013. arXiv:1302.4618.

[7] Heinz H. Bauschke, Patrick L. Combettes, and D. Russell Luke. Phase retrieval, error reduction algorithm, and Fienup variants: A view from convex optimization. *J. Opt. Soc. Amer. A*, 19:1334–1345, 2002.

[8] Bernhard G. Bodmann and Nathaniel Hammen. Stable phase retrieval with low-redundancy frames. *arXiv e-prints*, February 2013. arXiv:1302.5487.

[9] Jameson Cahill, Peter G. Casazza, Jesse Peterson, and Lindsey Woodland. Phase retrieval by projections. *arXiv e-prints*, May 2013. arXiv:1305.6226v3.

[10] Emmanuel J. Candès, Thomas Strohmer, and Vladislav Voroninski. PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.

[11] Aldo Conca, Dan Edidin, Milena Hering, and Cynthia Vinzant. An algebraic characterization of injectivity in phase retrieval. *arXiv e-prints*, December 2013. arXiv:1312.0158.

[12] Valentina Davidoiu, Bruno Sixou, Max Langer, and Françoise Peyrin. Nonlinear phase retrieval using projection operator and iterative wavelet thresholding. *IEEE Signal Process. Lett.*, 19(9):579 – 582, 2012.

[13] James R. Fienup. Phase retrieval algorithms: a comparison. *Applied Optics*, 21(15):2758–2769, 1982.

[14] Ralph W Gerchberg and W. Owen Saxton. A practical algorithm for the determination of the phase from image and diffraction plane pictures. *Optik*, 35(2):237–246, 1972.

[15] Alexander Grothendieck and Jean Dieudonné. Éléments de géométrie algébrique iv, deuxième partie. *Publ. Math. IHES*, 24, 1965.

[16] Alexander Grothendieck and Jean Dieudonné. Éléments de géométrie algébrique iv, troisième partie. *Publ. Math. IHES*, 28, 1966.

[17] Franz J. Király and Martin Ehler. The algebraic approach to phase retrieval and explicit inversion at the identifiability threshold. *arXiv e-prints*, February 2014. arXiv:1402.4053.

[18] Franz J. Király, Paul von Bünau, Jan Saputra Müller, Duncan Blythe, Frank Meinecke, and Klaus-Robert Müller. Regression for sets of polynomial equations. *JMLR Workshop and Conference Proceedings*, 22:628–637, 2012.

[19] David Mumford. *The Red Book of Varieties and Schemes*. Lecture Notes in Mathematics. Springer-Verlag Berlin Heidelberg, 1999.

[20] Andrew E. Yagle and Amy E. Bell. One- and two-dimensional minimum and nonminimum phase retrieval by solving linear systems of equations. *IEEE Trans. Signal Process.*, 47(11):2978–2989, 1999.