Spring 5-2021

Non-Gaussian Measurements of Coherent States of Light for Metrology and Communication

Matthew DiMario
University of New Mexico

Matthew Thomas DiMario

Candidate

Physics and Astronomy

Department

This dissertation is approved, and it is acceptable in quality and form for publication:

Approved by the Dissertation Committee:

Francisco Elohim Becerra-Chavez , Chairperson

Ivan Deutsch

Keith Lidke

Alberto Marino

Non-Gaussian Measurements of Coherent States of Light for Metrology and Communication

by

Matthew Thomas DiMario

B.S, Physics, University of Vermont, 2014
B.S, Mathematics, University of Vermont, 2014
M.S., Physics, University of New Mexico, 2017

DISSERTATION

Submitted in Partial Fulfillment of the
Requirements for the Degree of

Doctorate of Philosophy
Physics

The University of New Mexico

Albuquerque, New Mexico

May, 2021

Dedication

To Rachel, for the unending love and support...

Acknowledgments

I would not have made it to this point without the help and support of so many of my friends, family, and fellow scientists. In particular that of my advisor Prof. Elohim Becerra who allowed me to forge my own scientific path and provided much guidance over the course of my PhD. The past year (2020) in particular certainly has been a strange one, and I am grateful that things did not come to a screeching halt.

I also want to thank the dissertation committee: Profs. Ivan Deutsch, Alberto Marino, and Keith Lidke. This thesis would also have never been possible without the help of Gloria Cordova and Alisa Gibson, who kept me, CQuIC, and the Physics department on track on a regular basis. Thank you as well to everyone who is part of the CQuIC family, I could not have asked for a better scientific home.

I want to thank my mom for all her encouragement and for printing out my papers so that I know at least one person read them. I also want to thank my dad for being my 24/7 engineer-on-call when it comes to all things FPGA and brewing. Finally, a special thank you to Rachel for all of her love and patience over the years.

Non-Gaussian Measurements of Coherent States of Light for Metrology and Communication

by

Matthew Thomas DiMario

B.S, Physics, University of Vermont, 2014

B.S, Mathematics, University of Vermont, 2014

M.S., Physics, University of New Mexico, 2017

Ph.D., Physics, University of New Mexico, 2021

Abstract

Conventional measurement technology is unable to extract the most amount of information possible from coherent states of light. Non-Gaussian measurements which can count individual photons can surpass the sensitivity limits of ideal conventional strategies, and approach the ultimate limits achievable given by quantum mechanics. This thesis presents investigations and demonstrations of these unconventional measurements, which utilize coherent operations and single photon counting. This thesis shows that non-Gaussian measurements can outperform conventional strategies in estimation tasks as well as a variety of communication problems. This thesis also investigates novel approaches and algorithms for building robustness to static and dynamic noise which is present in realistic implementations, a critical barrier to transitioning non-Gaussian measurements out of the lab. Overall, this thesis aims to show that the single photon counting can be leveraged to implement measurements that operate at sensitivities which are inaccessible to current technology.

Contents

List of Figures

Chapter 1

Introduction

The process of measurement allows for probing the properties of a physical system and extracting information. The field of metrology seeks to answer the fundamental questions of how to perform the best possible measurement, and what are the ultimate limits in measurement precision [1, 2]. Generally, we aim to find strategies which have the smallest possible measurement uncertainty in their outcomes [3–5]. However, the fact that measurements of quantum systems are inherently *probabilistic* means that the same system may have many possible measurement outcomes, which limits the attainable uncertainty [6–9]. Nevertheless, the ultimate goal of metrology is to identify measurements which operate at the limits of sensitivity and uncertainty defined by the laws of quantum mechanics [7, 10–13].

Measurements which attempt to discriminate between a finite set of quantum states are a central aspect of classical and quantum information theory [2, 14–17]. In particular, discrimination of different states of laser light is a critical task in optical communication [18–23], quantum key distribution and cryptography [24–31], non-destructive measurements [32, 33], and even quantum computation [34–37]. On the other hand, a canonical problem in metrology is estimation of a single parameter, such as a phase shift acquired by a field of light after interacting with an object [4, 12, 38]. Phase measurements which achieve minimal estimation uncertainty have wide ranging

applications from biology [39–43] to magnetometry [44–46] to astronomy [12, 47], and quantum measurement theory provides a toolbox for identifying and constructing such measurements [8, 9, 48]. Regardless of whether information is encoded into just a few states or infinitely many, the unavoidable overlap between different states of laser light (i.e. coherent states) prevents any physical measurement from achieving perfect identification [49], even in principle [2, 7]. To this end, measurement theory gives a description of the ultimate sensitivity limits when identifying quantum states [2], and how to potentially reach these limits. Moreover, there is a significant gap between the sensitivity limits of current technology and what is fundamentally achievable [50].

A major goal of quantum measurement theory is finding measurement strategies which are physically realizable and surpass the sensitivity limits of current technology, or even approach the ultimate limits of state discrimination and estimation. On one hand, conventional measurements for phase estimation of laser light are ubiquitous, but strategies which approach the ultimate limits in estimation uncertainty for laser light have remained elusive. On the other hand, quantum measurements for discriminating multiple states of laser light have been successfully implemented in the laboratory, showing unconditional advantages over ideal conventional measurements [51–55]. However, a critical barrier to transitioning these measurements out of the lab and into practical applications is robustness to noise and imperfections in the communication channel and measurement itself. Noise can originate from the fact that the components of an experimental demonstration, such as beam splitters and detectors, are intrinsically imperfect. However, even with an ideal implementation there may be noise induced by a communication channel which will degrade the performance if not taken into account by the measurement. In any case, such noise and imperfections decrease performance in state discrimination and reduce the amount of decoded information. Solutions to these issues of robustness to noise in communication systems using conventional measurements currently enable many current communication technologies, albeit with performance far from the ultimate limits of information transfer. However, these approaches to combat the effects of noise and imperfections are tailored to conventional

measurements. Thus, realizations of quantum measurement strategies require novel methods to compensate for their own imperfections, as well as ways to maintain their performance in the presence of external noise sources such as intrinsic channel noise.

This thesis presents investigations and demonstrations of quantum measurements that achieve sensitivities beyond conventional limits of detection. We also study methods for making these quantum measurement strategies robust to realistic noise and imperfections. Experimentally, we implement such quantum measurements by combining coherent operations and interference with the ability to resolve individual numbers of photons. The measurement strategies investigated in this thesis can either be adaptive or non-adaptive, and possibly optimized *a priori* or in real-time to achieve the best possible performance. In addition, we also investigate methods to correct for noise present in realistic communication channels based on an intuitive model as well as machine learning techniques. These approaches can enable future, realistic implementations of quantum measurements for communication and estimation beyond the sensitivity limits of conventional technology. Simply put, the overall aim of the work in this thesis is to investigate the seemingly simple question of: how does one implement a *good* measurement?

1.1 Preliminaries: Coherent States of Light and Measurement Operators

Coherent states of light are the quantum mechanical description of laser light [8]. Having such a description facilitates the study of quantum measurements for state discrimination and metrology using light from a laser. To define the coherent states, we start with the quadrature operators $\hat{x}$ and $\hat{p}$ of the electromagnetic field, which do not commute: $[\hat{x}, \hat{p}] = i$ ($\hbar = 1$). We then define the standard creation and annihilation (ladder) operators $\hat{a}$ and $\hat{a}^\dagger$ as: $\hat{a} = (\hat{x} + i\hat{p})/\sqrt{2}$ and $\hat{a}^\dagger = (\hat{x} - i\hat{p})/\sqrt{2}$. These operators $\hat{a}$ and $\hat{a}^\dagger$ allow us to define eigenstates $|n\rangle$ of $\hat{n} = \hat{a}^\dagger\hat{a}$ as: $\hat{n}|n\rangle = n|n\rangle$, which correspond to quantum states of the electromagnetic field with exactly n photons,

known as number states, or Fock states [8]. The ladder operators have the effect of either raising $\hat{a}^\dagger|n\rangle = \sqrt{n+1}|n+1\rangle$ or lowing $\hat{a}|n\rangle = \sqrt{n}|n-1\rangle$ between different energy eigenstates. Given these operators, we define a coherent state of light $|\alpha\rangle$ as eigenstates of the annihilation operator: $\hat{a}|\alpha\rangle = \alpha|\alpha\rangle$, where a coherent state is parameterized by a complex number α. By expanding into the orthonormal basis of the Fock states and applying the annihilation operator, we obtain the relation:

$$\hat{a}|\alpha\rangle = \hat{a}\sum_{n=0}^{\infty} c_n|n\rangle = \sum_{n=1}^{\infty} c_n\sqrt{n}|n-1\rangle = \sum_{n=0}^{\infty} c_{n+1}\sqrt{n+1}|n\rangle = \alpha\sum_{n=0}^{\infty} c_n|n\rangle. \quad (1.1)$$

From the recursive relationship $\alpha c_n = c_{n+1}\sqrt{n+1}$ and normalization $\langle\alpha|\alpha\rangle = 1$, we arrive at a complete quantum mechanical description of a coherent state of light:

$$|\alpha\rangle = e^{-|\alpha|^2/2}\sum_{n=0}^{\infty}\frac{(\alpha\hat{a}^\dagger)^n}{n!}|0\rangle = e^{-|\alpha|^2/2}\sum_{n=0}^{\infty}\frac{\alpha^n}{\sqrt{n!}}|n\rangle, \quad (1.2)$$

where $|\alpha|^2 = \langle\alpha|\hat{n}|\alpha\rangle$ is referred to as the mean photon number of the state. Coherent states of light have many important properties including [56]:

- The probability $P(n)$ of detecting n photons from a coherent state of light is given by the Poisson distribution $P(n) = |\langle n|\alpha\rangle|^2 = \frac{(|\alpha|^2)^n}{n!}e^{-|\alpha|^2}$ with mean $|\alpha|^2$ and variance $\langle(\Delta\hat{n})^2\rangle = |\alpha|^2$.
- $\mathrm{Re}[\alpha] = \langle\hat{a}+\hat{a}^\dagger\rangle/2 = \langle\alpha|\hat{x}|\alpha\rangle/\sqrt{2}$ and $\mathrm{Im}[\alpha] = \langle\hat{a}-\hat{a}^\dagger\rangle/2i = \langle\alpha|\hat{p}|\alpha\rangle/\sqrt{2}$.
- Like the vacuum state $|0\rangle$, $|\alpha\rangle$ is a *minimum uncertainty* state in that $\langle(\Delta\hat{x})^2\rangle = \langle(\Delta\hat{p})^2\rangle = 1/2$ and the marginal quadrature distributions $|\langle x|\alpha\rangle|^2 = \frac{1}{\sqrt{\pi}}e^{-(x-\langle\hat{x}\rangle)^2}$ and $|\langle p|\alpha\rangle|^2 = \frac{1}{\sqrt{\pi}}e^{-(p-\langle\hat{p}\rangle)^2}$ are Gaussian distributed.
- Coherent states are *not* orthonormal: $\langle\alpha|\beta\rangle = e^{-(|\alpha|^2+|\beta|^2-2\alpha^*\beta)/2}$. However, they form an overcomplete basis satisfying $\hat{I} = \frac{1}{\pi}\int d^2\alpha|\alpha\rangle\langle\alpha|$.

Figure 1.1(a) shows the phase space representation of a coherent state of light $|\alpha\rangle$. The blue region shows the fundamental quantum noise associated with coherent states and this fundamental uncertainty satisfies Heisenbergs uncertainty relation,

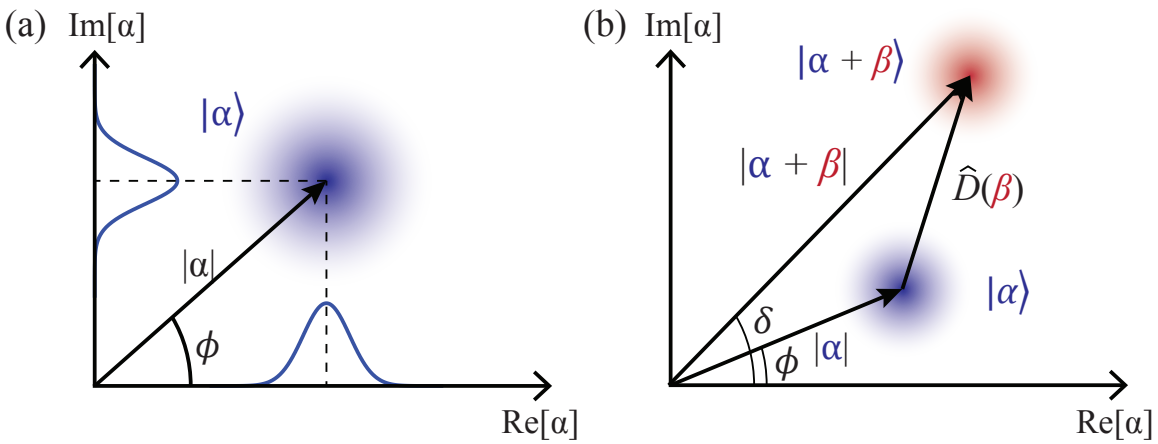

Figure 1.1: (a) A coherent state if light $|\alpha\rangle$ in phase space. The blue region shows the intrinsic quantum noise associated with coherent states, which satisfies Heisenbergs uncertainty relation. (b) Phase space picture of a displacement operation $\hat{D}(\beta)$ acting on $|\alpha\rangle$ to obtain the state $\hat{D}(\beta)|\alpha\rangle = |\alpha + \beta\rangle$.

i.e. $\langle(\Delta\hat{x})^2\rangle\langle(\Delta\hat{p})^2\rangle = 1/4$ with $\hbar = 1$. Coherent states of light can alternatively be defined with respect to the displacement operator $\hat{D}(\alpha) \equiv e^{\alpha\hat{a}^\dagger - \alpha^*\hat{a}}$, which has two convenient properties [8]:

- $\hat{D}^\dagger(\alpha) = \hat{D}(-\alpha)$.
- $\hat{D}(\alpha)\hat{D}(\beta) = e^{\alpha\beta^* - \alpha^*\beta}\hat{D}(\alpha + \beta)$.

Thus, we can equivalently define a coherent state as the displacement operator acting on the vacuum: $|\alpha\rangle \equiv \hat{D}(\alpha)|0\rangle$. In addition, the displacement operator can act on a coherent state such that $\hat{D}(\beta)|\alpha\rangle = |\alpha + \beta\rangle$, where by adjusting the value of β, we can translate $|\alpha\rangle$ into a different coherent state which can have a different photon number distribution. Figure 1.1(b) shows the phase space picture of a displacement operation $\hat{D}(\beta)$ acting on the state $|\alpha\rangle$ to obtain the state $\hat{D}(\beta)|\alpha\rangle = |\alpha + \beta\rangle$.

The probability of detecting n photons from a coherent state $|\alpha\rangle$ is given by $P(n) = \mathrm{Tr}\big[|n\rangle\langle n||\alpha\rangle\langle\alpha|\big] = |\langle n|\alpha\rangle|^2$, which is obtained through the typical picture of a projective measurement $|n\rangle\langle n|$ onto the Fock basis. However, measurements fundamentally need not be projective. To describe a general measurement, one can

define a positive operator valued measure (POVM) which is a set of positive-semi-definite operators $\{\hat{\Pi}_i\}$ on a Hilbert space which sum to the identity: $\sum_i \hat{\Pi}_i = \hat{I}$ [9, 57]. Each POVM *element* $\hat{\Pi}_i$ corresponds to a particular measurement setting such that if a general quantum state described by the density matrix $\hat{\rho}$ is measured, the probability of outcome i is: $P(i) = \mathrm{Tr}[\hat{\Pi}_i\hat{\rho}]$. In the case where $\hat{\Pi}_i = |\psi_i\rangle\langle\psi_i|$ is a projector and $\hat{\rho} = |\phi\rangle\langle\phi|$ is a pure state, the traditional picture is recovered as $P(i) = \mathrm{Tr}[|\psi_i\rangle\langle\psi_i||\phi\rangle\langle\phi|] = |\langle\psi_i|\phi\rangle|^2$.

The POVM language of measurement theory allows for constructing more general classes of measurements which go beyond simple projective measurements. As discussed in Chapter 2, this description of measurements becomes beneficial when comparing measurement strategies that are designed to optimize different physical quantities, such as minimize measurement errors [2], maximize information transfer [58, 59], or minimize inconclusive results [15]. This POVM framework also allows for constructing measurements which may operate on multiple modes simultaneously, or in a non-local way [8, 9, 57]. As discussed throughout this thesis, measurements are an extremely critical aspect of many communication and metrological tasks. The fundamental limits for these tasks are often known [2], but the POVM which reaches that limit may not be physically realizable in the laboratory. This division between theory and experiment leads to the question of which POVM is experimentally implementable and achieves the best possible performance.

1.2 Probability of Error in Discrimination

A receiver implementing a measurement strategy for state discrimination knows *a priori* that the input state is drawn from a finite set of possible states. Thus, the input state is decoded either correctly or incorrectly given a particular measurement outcome. Of course, the receiver is ideally designed to provide the correct answer to the state discrimination problem as often as possible, i.e. have a low probability of error P_E and high probability of correct detection P_C. For an arbitrary set of M

input states $\{\hat{\rho}_k\}_M$ with prior probabilities $\{p_k\}_M$, the most general expression for the probability of error P_E is:

$$P_E = 1 - P_C = 1 - \sum_{k=1}^{M} p_k P(\hat{\rho}_k|\hat{\rho}_k) = 1 - \sum_{k=1}^{M} p_k \mathrm{Tr}\big[\hat{\Pi}_k \hat{\rho}_k\big]. \tag{1.3}$$

The quantity $P(\hat{\rho}_k|\hat{\rho}_k) = \mathrm{Tr}\big[\hat{\Pi}_k\hat{\rho}_k\big]$ is the probability of guessing state $\hat{\rho}_k$, which corresponds to the measurement outcome $\hat{\Pi}_k$, given that the input state was $\hat{\rho}_k$. The guess for the input state in discrimination problems usually corresponds to the most likely state given the particular measurement outcome, i.e. the state with the highest *posterior* probability. Typically, thought not always, the main goal of the problem of state discrimination is to find strategies for which minimize P_E for a given set of states [15]. This task translates into finding an optimal POVM, known as the Helstrom measurement, which has the smallest possible error probability, referred to as the Helstrom bound [2, 14].

The Helstrom measurement for binary pure states $\{\hat{\rho}_1, \hat{\rho}_2\} = \{|\psi_1\rangle\langle\psi_1|, |\psi_2\rangle\langle\psi_2|\}$ is described by $\{\hat{\Pi}_1, \hat{\Pi}_2\}$ where $\hat{\Pi}_{1,2}$ are projectors onto the eigenvectors of $p_1\hat{\rho}_1 - p_2\hat{\rho}_2$. A positive outcome of $\hat{\Pi}_1$ indicates that $\hat{\rho}_1$ is the more likely input state, and vice versa. The Helstrom bound on the error probability is then given by:

$$P_E^{Hel} = \frac{1}{2}\Big(1 - \sqrt{1 - 4p_1p_2|\langle\psi_1|\psi_2\rangle|^2}\Big). \tag{1.4}$$

The Helstrom measurement for discriminating $\{|\psi_1\rangle, |\psi_2\rangle\}$ when $p_{1,2} = 1/2$ is implemented by projecting onto the orthonormal basis states $|\phi_\pm\rangle$, where $|\psi_1\rangle = a|\phi_+\rangle + b|\phi_-\rangle$ and $|\psi_2\rangle = b|\phi_+\rangle + a|\phi_-\rangle$. Figure 1.2 shows the vector representation of the two states $|\psi_{1,2}\rangle$ (black) and the Helstrom measurement $|\phi_\pm\rangle$ (red), where the angle θ characterizes the non-orthogonality between $|\psi_1\rangle$ and $|\psi_2\rangle$.

In the specific case of two coherent states with opposite phase $\{|\alpha\rangle, |-\alpha\rangle\}$ and equal prior probabilities, the Helstrom measurement corresponds to projecting onto the so-called cat states: $N_\pm(|\alpha\rangle \pm |-\alpha\rangle)$ [2]. While effort has been made into how to *approximately* implement these kinds of measurements [60], and make them non-destructively [33, 61], the Helstrom measurement for coherent states has yet to be

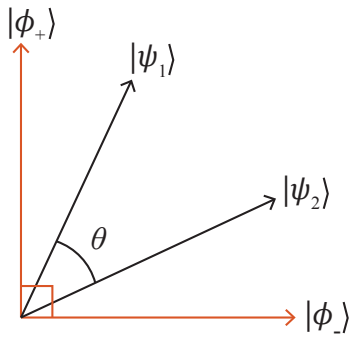

Figure 1.2: Vector representation of the binary states $|\psi_1\rangle$ and $|\psi_2\rangle$ (black) and the Helstrom measurement (red), which is formed by projecting onto the basis states $|\phi_+\rangle$ and $|\phi_-\rangle$.

demonstrated. Of note however, is that for binary coherent states there exists a measurement strategy which can reach the Helstrom bound using readily available technology. This strategy is known as the Dolinar receiver and is based on displacement operations, single photon detection, and infinite bandwidth feedback [53, 62–64], and we study this measurement in Chapter 2.

In contrast to the problem of binary state discrimination, the optimal measurement for the problem of discriminating many non-orthogonal states is less intuitive, if known at all [15]. A particularly important class of states which exhibit a high degree of symmetry is when every state $\hat{\rho}_k = |\psi_k\rangle\langle\psi_k|$ is generated by repeatedly applying the same unitary operation $\hat{U}$: $|\psi_k\rangle = \hat{U}|\psi_{k-1}\rangle = (\hat{U})^k|\psi_0\rangle$, and $p_k = 1/M$, known as geometrically uniform states [15, 65]. For example, a particular state in the four state alphabet $\{|\alpha e^{ik\pi/2}\rangle\}$ with $k \in \{0, 1, 2, 3\}$ is obtained by applying the rotation operator $\hat{U} = e^{i\hat{n}\frac{\pi}{2}}$ to the state $|\alpha\rangle$ ($k = 0$) a particular number of times. The optimal POVM for these states is the square-root measurement and given by $\hat{\Pi}_k = \frac{1}{M}\sqrt{\hat{\rho}}\hat{\rho}_k\sqrt{\hat{\rho}}$ where $\hat{\rho} = \frac{1}{M}\sum_k \hat{\rho}_k$, often corresponding to projections onto complex superposition states [65–67]. Thus, the Helstrom bound for this special but ubiquitous class of states is given by [59]:

$$P_E^{Hel} = 1 - \frac{1}{M^2}\left(\sum_{k=1}^{M}\sqrt{\lambda_k}\right)^2, \tag{1.5}$$

where λ_k is the k^{th} eigenvalue of the Gram matrix $\mathbf{G}$ with $G_{i,j} = \langle\psi_i|\psi_j\rangle$.

A closed form expression for the optimal POVM $\{\hat{\Pi}_k\}$ for state discrimination in general cannot be found for an arbitrary set of M states $\{\hat{\rho}_k\}_M$ with prior probabilities $\{p_k\}_M$ [15]. Furthermore, the physical implementations of these optimal measurements for coherent states of light are unknown. Thus, physically realizable measurement strategies which approach the ultimate bounds of discrimination are not only of fundamental importance, but also practically relevant. This is especially true when considering applications of state discrimination to optical communication where the limits of information transfer are set by how well coherent states of light can be distinguished [50, 58, 68]. Thus, to efficiently facilitate the transfer of information, one needs to account for how information is encoded into coherent states of light as well as possible measurement strategies [18, 58, 59].

1.3 Encoding Information into Physical States

Considering a particular communication scenario, a sender and receiver need to implement a certain encoding of information into a physical set of coherent states, generally with a restriction on the *average* energy of the states [58]. In the context of communication with binary states, a scheme which encodes information into the magnitude of two coherent states is referred to as on-off-keying (OOK) and corresponds to the alphabet $|\alpha_k\rangle = \{|0\rangle, |\alpha\rangle\}$. However, for a given amount of average energy $\bar{n} = \frac{1}{2}(|\alpha_1|^2 + |\alpha_2|^2) = |\alpha|^2$ between the two coherent states, an OOK alphabet ($\{|0\rangle, |\sqrt{2}\alpha\rangle\}$) is not optimal. We refer to an optimal encoding for error probability as one which maximizes the distinguishability of the states within the alphabet under an average energy constraint. For the purposes of this thesis, distinguishability is equivalent to the amount of overlap, or non-orthogonality, between the states in the alphabet. Thus, the distinguishability of binary coherent states given the average energy $\bar{n}$ is $|\langle\alpha_1|\alpha_2\rangle|^2$. Encoding information into the phase of the light to obtain a binary phase-shift-keyed (BPSK) alphabet: $|\alpha_k\rangle = \{|-\alpha\rangle, |\alpha\rangle\}$ maximizes the energy efficiency of a binary alphabet [69]. Here α is assumed to be real and positive and the

information is encoded into the phase as either 0 or π. It is important to note that if a coherent encoding is used, then the receiver must employ a coherent detection scheme [50, 58].

Regardless of the encoding, a communication scheme which utilizes two states can only transmit at most a single bit of information. The sender can employ a larger alphabet consisting of M states to increase the possible amount of information transferred through a communication channel to $\log_2(M)$. A natural extension is from a single bit to two bits of information, which requires an alphabet of $M = 4$ coherent states. The optimal encoding in terms of energy efficiency is now a quaternary phase-shift-keyed (QPSK) alphabet: $|\alpha_k\rangle \in \{|\alpha\rangle, |i\alpha\rangle, |-\alpha\rangle, |-i\alpha\rangle\}$ where the information is solely in the phase of the states. To further increase the amount of information encoded to four bits using $M = 16$ states, one may expect that the optimal alphabet requires coherent states with 16 different phases. However, given that the effectiveness of an encoding scheme is with respect to the *average* energy, a more efficient alphabet is one which leverages the ability to encode information into the phase and amplitude simultaneously [58, 68]. These encodings are referred to as quadrature-amplitude modulation (QAM) where states are usually arranged in a grid-like "constellation" in phase space, i.e. 16-QAM corresponds to a 4×4 grid of states. Figure 1.3(a-c) shows examples of different encoding schemes which employ 2, 4, and 16 states, respectively. Although more information can be encoded into higher order constellations, the error in discrimination also increases. Figure 1.3(d) shows the how the Helstrom bound for the different encodings in (a-c) compare as a function of the mean photon number $\bar{n}$, where $\bar{n} = \frac{1}{M}\sum_k |\alpha_k|^2$. The increase in discrimination error is due to the fact that with more states, there is more overlap and thus a higher degree of non-orthogonality between the states in the alphabet [70].

A communication strategy can also aim to be optimal in terms of information efficiency, as opposed to error probability. In such a scenario, the strategy seeks to maximize the amount information transfer for a given average energy. We can define the average energy more generally as $\bar{n} = \sum_k p_k |\alpha_k|^2$ where p_k is the prior probability of

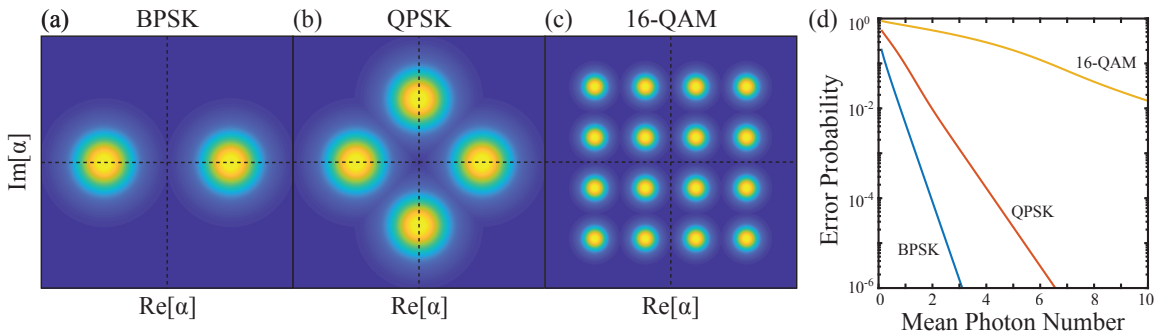

Figure 1.3: (a-c) Examples of the optimal coherent state alphabets shown in phase space when using (a) 2 coherent states, (b) 4 coherent state, and (c) 16 coherent states. (d) Helstrom bound on the error probability for each alphabet in (a-c) as a function of the average mean photon number.

the transmitted state. Communication strategies which achieve the highest information efficiency optimize both the prior probabilities $\{p_k\}_M$ as well as the physical states $\{|\alpha_k\rangle\}_M$. The optimal scheme for a pure loss channel imposes a Gaussian prior probability distribution on the states within the alphabet as $p_k \propto e^{-|\alpha_k|^2/2\sigma^2}$, and σ and $|\alpha_k|^2$ are optimized for a given $\bar{n}$ [18]. We study how non-Gaussian measurements can increase the amount of information transfer when combined with non-uniform prior probability distributions for 16-QAM encodings in Ch. 6. A pure loss channel is an example of the well studied class of Gaussian communication channels, where the ultimate limits of information transfer are known [18, 71]. On the other hand, if the communication channel imparts non-Gaussian noise onto the states then the ultimate limits of information transfer may be unknown [18, 72–75], and encodings which were optimal for a loss channel may not be ideal in the presence of non-trivial channel noise. We study the problem of optimal communication strategies over non-Gaussian channels in Ch. 2.2 in the context of binary communication.

1.4 Defining the Quantum Noise Limit

The fundamental limit of the Helstrom bound and how information is encoded is only half of the story. The information needs to be decoded through a physically realizable

measurement of the states of light once the states propagate through a communication channel. Current measurement technologies for coherent optical communication mainly rely on strategies which implement Gaussian measurements of the input state [5, 76]. A Gaussian measurement is one with outcomes which exhibit Gaussian distributed statistics. The formalism of Gaussian measurement theory [5, 8, 9, 57, 76] allows us to calculate the fundamental limit of the error probability for Gaussian measurements [5]. This limit of Gaussian receivers is referred to as the quantum noise limit (QNL) and in general is much worse than the Helstrom bound. Thus, the QNL serves as the benchmark of conventional measurements.[1] Our investigations focus on studying unconventional non-Gaussian measurements which unconditionally (without correcting for any inefficiencies) surpass the QNL and push towards the Helstrom bound.

The ideal Gaussian measurement for the BPSK alphabet $\{|\pm\alpha\rangle\}$ is an optical homodyne receiver, which measures the x-quadrature (real component) of the state by interfering the input field with a strong local oscillator field with phase $\pi/2$, followed by balanced intensity measurements [50, 58]. This measurement is described by the POVM $\{|x\rangle\langle x|\}$ and the probability of correct detection P_C for a homodyne receiver corresponds to [69, 77]:

$$P_C(\pm\alpha|\pm\alpha) = \frac{1}{\sqrt{\pi}}\int\limits_{R_\pm} \mathrm{Tr}\Big[|\pm\alpha\rangle\langle\pm\alpha||x\rangle\langle x|\Big]dx = \frac{1}{\sqrt{\pi}}\int\limits_{R_\pm} e^{-(x\mp\sqrt{2}\alpha)^2}dx \tag{1.6}$$

$$= \Big[1+\mathrm{erf}(\sqrt{2}\alpha)\Big], \tag{1.7}$$

where $\mathrm{erf}(\cdot)$ is the error function. The integration region $R_\pm$ is the region of phase space where $|\pm\alpha\rangle$ is the most likely state given the measurement outcome x. This integration is performed only over certain regions instead of the entire phase space due to the fact that the guess for the input state is the most likely state given the measurement outcome, i.e. the state with the largest posterior probability. Thus, the

[1] We note that this QNL for state discrimination is defined with respect to conventional Gaussian measurements and assuming the use of coherent states of light. This is in contrast to the so-called standard quantum limit for estimation and sensing tasks, which is a statement about the achievable uncertainty when assuming just a coherent state input and is discussed in the context of phase estimation in Chapter 5.

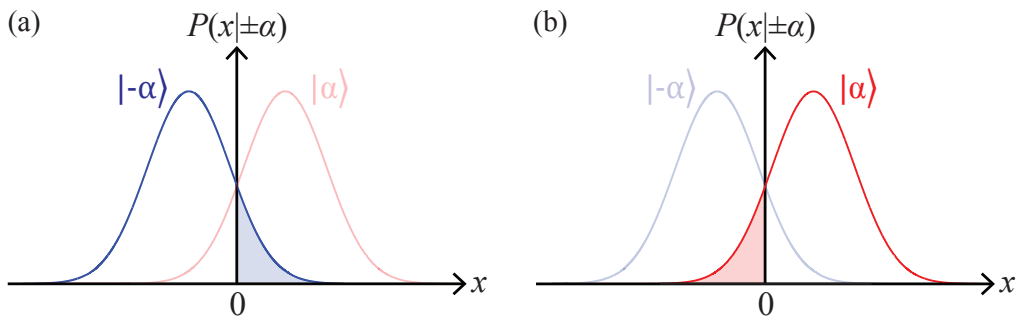

Figure 1.4: Likelihood functions $P(x|\pm\alpha)$ for the optimal homodyne measurement of the binary coherent states (a) $|-\alpha\rangle$ and (b) $|\alpha\rangle$. The blue and red shaded regions correspond to measurement outcomes which result in discrimination errors.

total probability of correctly detecting a particular input state is the combination of only the detection results for which that input state has the largest posterior probability. In fact, the homodyne receiver represents the optimal Gaussian measurement for BPSK states [69] and serves as the QNL for this problem:

$$\mathrm{QNL}_{BPSK} = 1 - P_C = 1 - \left(\frac{1}{2}P(\alpha|\alpha) + \frac{1}{2}P(-\alpha|-\alpha)\right) \tag{1.8}$$

$$= 1 - \frac{1}{2}\Big[1 + \mathrm{erf}(\sqrt{2}\alpha)\Big] = \frac{1}{2}\Big[1 - \mathrm{erf}(\sqrt{2}\alpha)\Big]. \tag{1.9}$$

Figure 1.4(a,b) shows the likelihood function $P(x|\pm\alpha)$ of each detection outcome x for the two coherent states $|\pm\alpha\rangle$. Positive outcomes ($x \geq 0$) correspond to guessing the input state $|\alpha\rangle$ and negative outcomes ($x < 0$) correspond to $|-\alpha\rangle$. The total probability $P(\alpha|\alpha)$ of correctly detecting $|\alpha\rangle$ is obtained from integrating $P(x|\alpha)$ from $x = 0$ to $x = \infty$ (i.e. $R_+ = [0, \infty]$), and vise versa for $|-\alpha\rangle$. Discrimination errors occur when either a positive outcome is obtained from $|-\alpha\rangle$ (shaded region in (a)), or when a negative outcome is obtained from $|\alpha\rangle$ (shaded region in (b)).

Similarly, the conventional Gaussian receiver for a QPSK alphabet is an ideal heterodyne measurement, or double-homodyne, which splits the input state in two modes and measures the x-quadrature of one mode, and the p-quadrature of the other [8, 58]. Given a measurement outcome γ, which is a complex number, this can be seen as projecting onto a coherent state $|\gamma\rangle\langle\gamma|$ such that heterodyne detection is described by the POVM $\{\frac{1}{\pi}|\gamma\rangle\langle\gamma|\}$ [58, 59, 76]. The probability of correct detection P_C

is the product of P_C for two homodyne measurements such that the QNL for QPSK states is given by:

$$\mathrm{QNL}_{QPSK} = 1 - \frac{1}{4}\left[1 + \mathrm{erf}\left(\frac{\alpha}{\sqrt{2}}\right)\right]^2. \tag{1.10}$$

For an arbitrary coherent state alphabet $\{|\alpha_k\rangle, p_k\}_M$ which spreads over both phase space dimensions, the limit of an ideal heterodyne measurement serves as the general quantum noise limit [59]:

$$\mathrm{QNL} = 1 - \sum_{k=1}^{M} \frac{p_k}{\pi} \int_{R_k} |\langle \alpha_k | \gamma \rangle|^2 d^2\gamma, \tag{1.11}$$

where in general R_k is the phase space region where the state $|\alpha_k\rangle$ is the most likely state given the measurement outcome γ.

The root of many of the problems addressed in this thesis however, is that in general for the problem of state discrimination there is an exponentially large gap between the QNL and Helstrom bound [59]. Thus, implementable and imperfect strategies which can surpass the QNL and approach the ultimate limits of discrimination are of critical importance to optical communication, and more generally quantum information and measurement theory. Furthermore, conventional Gaussian measurements also do not reach the fundamental limits associated with metrological problems such as optical phase estimation of coherent states [78]. However, the limit of conventional strategies for phase estimation is closer to the fundamental limit than for the state discrimination problem. Therefore, non-Gaussian strategies must be incredibly well understood in order to surpass the limit of conventional measurements, and we investigate this problem of non-Gaussian receivers for optical phase estimation in Ch. 5.

1.5 Photon Counting Measurements

The fundamental measurement aspect which enables unconventional receivers to surpass the limits of conventional measurements is having a non-Gaussian element

such as photon counting. Instead of measuring the intensity of the light via the operator $\hat{a}^\dagger\hat{a}$, as with homodyne or heterodyne measurements, a receiver can measure individual photons by projecting onto individual Fock states, which is described by the POVM $\{|n\rangle\langle n|\}$. By itself, $\{|n\rangle\langle n|\}$ is not a coherent measurement because this POVM alone cannot decode phase information. Thus, the receiver needs to implement a coherent operation *before* photon counting. Homodyne and heterodyne measurements rely on interference on a 50/50 beam splitter between the input state and an infinitely (compared to the input) strong reference field of light, referred to as a local oscillator (LO) [8, 58, 68]. In contrast, we focus on using a coherent displacement operation $\hat{D}(\beta)$ implemented using a LO field with a precise intensity and a highly transmissive ($T \approx 1$) beam splitter [79]. We choose the state of the LO to be $|\beta/\sqrt{1-T}\rangle$ and for experimental convenience, we include an implicit π phase shift into the displacement such that $\hat{D}(\beta)|\alpha\rangle = |\alpha - \beta\rangle$. The displaced photon counting measurement is then described by the POVM $\{\hat{\Pi}_n(\beta) = \hat{D}(\beta)|n\rangle\langle n|\hat{D}^\dagger(\beta)\}$.

Given a particular measurement outcome corresponding to a single POVM element $\hat{\Pi}_n(\beta)$, the receiver can calculate the likelihood of that detection for each input state as well as the posterior probability using Bayes rule. The decision for the input state in the problem of state discrimination is given by the state which has maximum *a posteriori* probability. However, the receiver can also perform an adaptive strategy instead of measuring the entire input state at once. In adaptive schemes, the input is split into multiple adaptive steps, each of which implements a displaced photon counting measurement and the probability of each input state is updated by recursively applying Bayes rule (See Appendix A.1). Figure 1.5 shows a diagram of an adaptive photon counting measurement. We use an LO to displace the input in phase space using interference on a 99/1 beam splitter, followed by detection with a single photon detector (SPD). Adaptive measurement strategies incorporate feedback to the LO conditioned on the particular photon detection result.

While we may be able to write down the POVM for a single adaptive step as $\hat{\Pi}_n(\beta)$, constructing the POVM which describes the *entire* adaptive measurement is

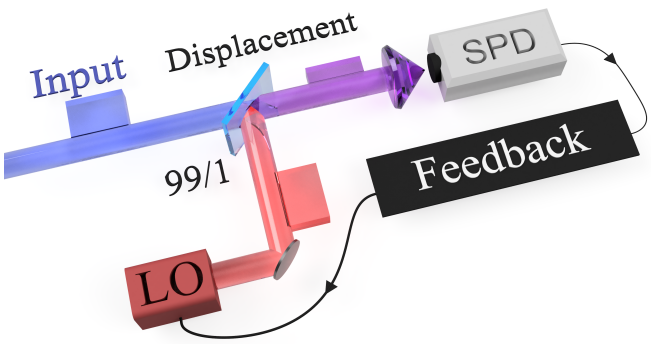

Figure 1.5: Diagram of an adaptive photon counting measurement. The input state (blue) is displaced in phase space using interference on a 99/1 beam splitter between the input and a local oscillator (LO). The displaced state is detected by a single photon detector (SPD). An adaptive strategy includes feedback to the LO based on the photon detection result.

more complex. Nevertheless, we aim to demonstrate that displaced photon counting measurements both adaptive and non-adaptive can surpass the limits of conventional technologies (i.e. Gaussian measurements) for a wide variety of problems. These investigations advance our understanding of how to best extract information encoded into coherent states of light, and further open the door to utilizing non-Gaussian measurements in quantum information science.

1.6 Thesis Organization

In this thesis, we investigate three main problems: (1) the discrimination of coherent states with different measurement strategies and their applications in communication, (2) methods for providing robustness for quantum measurements of coherent states to noise and imperfections in realistic communication channels and devices, and (3) the problem of phase estimation of coherent states with single-shot measurements. Chapter 2 discusses the problem of binary state discrimination in the context of communications and generalized quantum measurements. Section 2.1 describes how incorporating photon number resolution into an optimized non-adaptive strategy allows for robustness to experimental imperfections [80]. Section 2.2 investigates

using this robust receiver for a phase diffusion channel [81]. Section 2.3 presents the demonstration of optimally inconclusive measurement receivers for binary coherent states [64].

Chapter 3 addresses problems in the discrimination of QPSK coherent states based on optimized measurement strategies as well as strategies based on machine learning. Section 3.1 presents an optimized non-adaptive strategy based on a polarization interferometer [82]. Section 3.2 describes how different global optimizations for adaptive strategies can allow for unconditionally surpassing the QNL in the single-photon regime [83]. Section 3.3 investigates a reinforcement learning framework for adaptive photon counting measurements for state discrimination [84].

Chapter 4 investigates two methods for real-time noise tracking and correction for adaptive photon counting receivers in the presence of dynamic channel noise. Section 4.1 describes an experimental implementation of a phase tracking method based on a simple and practical estimator [85]. Section 4.2 presents a theoretical investigation of a scalable approach where a neural network is trained as a multi-parameter estimator both phase and amplitude noise [86].

Chapter 5 presents an adaptive photon counting measurement for optical phase estimation based on real-time optimizations [87, 88]. Section 5.2 describes the phase estimation strategy and Section 5.3 describes the experimental implementation. Section 5.4 presents the experimental and theoretical results. Section 5.5 discusses the approximations used to demonstrate the strategy, and the asymptotic connection between different optimization strategies and the Fisher information.

Chapter 6 discusses the possibility of using higher order modulation formats combined with an optimized distribution for the input state alphabet to maximize the mutual information. We find that adaptive photon counting measurements can reach rates of information transfer beyond the limit of conventional Gaussian strategies [89]. Chapter 7 summarizes the work and provides concluding remarks and future directions for non-Gaussian measurements. The published works that this thesis is based on and

their corresponding sections are:

- Chapter 2.1: **M. T. DiMario** and F. E. Becerra, "Robust Measurement for the Discrimination of Binary Coherent States", *Phys. Rev. Lett.* **121**, 023603 (2018).
- Chapter 2.2: **M. T. DiMario**, L. Kunz, K. Banaszek, and F. E. Becerra, "Optimized Communication Strategies with Binary Coherent States Over Phase Noise Channels", *npj Quantum Inf.* **5**, 65 (2019).
- Chapter 2.3: **M. T. DiMario** and F. E. Becerra, "Optimally Inconclusive Measurements of Binary Coherent States", *in preparation* (2021).
- Chapter 3.1: **M. T. DiMario**, E. C., R. A. J., and F. E. Becerra, "Implementation of a Single-Shot Receiver for Quaternary Phase-Shift Keyed Coherent States", *JOSA B* **35**, 568 (2018).
- Chapter 3.2: A. R. Ferdinand, **M. T. DiMario**, and F. E. Becerra, "Multi-State Discrimination Below the Quantum Noise Limit at the Single-Photon Level", *npj Quantum Inf.* **3**, 43 (2017).
- Chapter 4.1: **M. T. DiMario** and F. E. Becerra, "Phase Tracking for sub-Shot-Noise-Limited Receivers", *Phys. Rev. Research* **2**, 023384 (2020).
- Chapter 4.2: **M. T. DiMario** and F. E. Becerra, "Channel Noise Tracking for sub-Shot-Noise-Limited Receivers with Neural Networks", *Phys. Rev. Research* **3**, 013200 (2021).
- Chapter 5: **M. T. DiMario** and F. E. Becerra, "Single-Shot Non-Gaussian Measurements for Optical Phase Estimation", *Phys. Rev. Lett.* **125**, 120505 (2020).

Chapter 2

Binary Encoded Coherent States

The Helstrom bound is the smallest possible error probability in discriminating the binary coherent states $|\alpha_k\rangle = \{|-\alpha\rangle, |\alpha\rangle\}$. The Helstrom measurement is made by projecting onto the superposition state $N_{\pm}(|\alpha\rangle \pm |-\alpha\rangle)$, but has no known physical implementation. However, this ultimate bound can be reached by a measurement known as a Dolinar receiver, which has a relatively complex implementation that is susceptible to noise [53, 62, 90]. On the other hand, simple and robust receivers which can surpass the QNL rely on a fixed displacement operation $\hat{D}(\beta)$ (β real and positive) followed by single-photon detection. The first receiver of this kind was initially described by Kennedy [91], where $|\beta| = |\alpha|$ such that the receiver attempts to displace one of the states to the vacuum state followed by PNR(1) (on/off) detection, giving an error probability of: $P_E = e^{-4|\alpha|^2}/2$ for equiprobable states. This measurement was further improved to the so-called optimized Kennedy receiver by optimizing the displacement magnitude $|\beta| = |\beta_{opt}|$ such that the overall probability of error is minimized [69]. This optimized Kennedy receiver can in principle outperform ideal homodyne detection, the optimal Gaussian measurement, for all input powers. Experimentally, the optimized Kennedy strategy has been demonstrated surpassing the QNL in the low power limit ($|\alpha|^2 \leq 1$) both with [92] and without [52] correction

for the system efficiency, corresponding to unconditional performance below the QNL.

While the demonstrations in Refs. [92] and [52] represent a significant advance in understanding the benefit and capabilities of non-Gaussian measurements, the implementations are susceptible to experimental imperfections such as non-ideal interference visibility or detector dark counts. These imperfections limit the benefit of the optimized Kennedy receiver over the QNL to the low power regime if the measurement cannot be made robust to noise [80]. In addition, realistic communication channels induce noise in the coherent states, such as thermal noise or phase noise as opposed to a pure loss channel [93–95]. Such channel noise will impact the ability of a receiver to perform state discrimination due to the noise reducing the distinguishability of the states. A robust communication strategy also optimizes the input encoding because a BPSK alphabet may not provide the best distinguishability [81] after propagation through a noisy communication channel. Furthermore, the optimized Kennedy receiver and similar approaches fall into the paradigm of minimum error state discrimination, where the aim is to achieve the smallest possible error probability for a fixed energy of the input states. However, there exist other measurement strategies which aim to be unambiguous in their outcomes such that $P_E = 0$ [96, 97] at the cost of having inconclusive results. Moreover, strategies have been developed [17, 63] which interpolate between a minimum error and unambiguous receivers to implement optimally inconclusive measurements [64], a more general class of measurement strategies.

In this chapter, we investigate how to increase the robustness of non-Gaussian measurements of binary coherent states, and extend the state discrimination toolbox to include more general measurements beyond minimum error discrimination. In Section 2.1, we leverage the ability to resolve individual numbers of photons, as opposed to on/off detection, for BPSK states to provide the measurement with information about imperfections and enable robust discrimination strategies [80]. In Section 2.2, we study optimized measurements for binary state discrimination in the presence of phase noise channels [81], for which the ultimate limits of communication are not known. In

Section 2.3, we investigate how to implement optimally inconclusive measurements [64], which can achieve sub-Helstrom error probabilities at the cost of sometimes producing an inconclusive result.

2.1 Robust Optimized Kennedy with PNR

An ideal optimized Kennedy receiver for BPSK coherent states is able to surpass the QNL for all input mean photon numbers. However, this performance is degraded due to noise and imperfections in the measurement, and previous sub-QNL demonstrations [52, 92] have thus been limited to the low mean photon number regime ($|\alpha|^2 \leq 1$). A missing component which allows for surpassing the QNL beyond this regime is robustness to experimental imperfections such as non-ideal interference visibility $\xi \neq 1$ and detector dark counts $\nu \neq 0$. The overall effect of these imperfections is an increase in the error probability relative to the QNL as the mean photon number of the states increases due to erroneous photon detections. To counteract experimental imperfections, the receiver can leverage photon number resolving capabilities to recover sub-QNL performance by factoring in the information gained by detecting a precise number of photons. This photon number resolution allows the receiver to incorporate imperfections when optimizing the strategy as well as making an estimate for the input state by adjusting which measurement outcomes correspond to each input state.

2.1.1 Measurement Strategy

Figure 2.1(a) shows the strategy for the robust optimized Kennedy receiver with photon number resolution (PNR). The input state $|\alpha_k\rangle \in \{|-\alpha\rangle, |\alpha\rangle\}$ is first displaced in phase space by $\hat{D}(\beta)$ using interference with a local oscillator (LO) on a highly-transmissive beam splitter [79]. The displaced state $\hat{D}(\beta)|\pm\alpha\rangle = |\pm\alpha - \beta\rangle$ is then detected by a single photon detector (SPD) with PNR(m), where m is the number of photons the detector can resolve before becoming a threshold detector. We note that

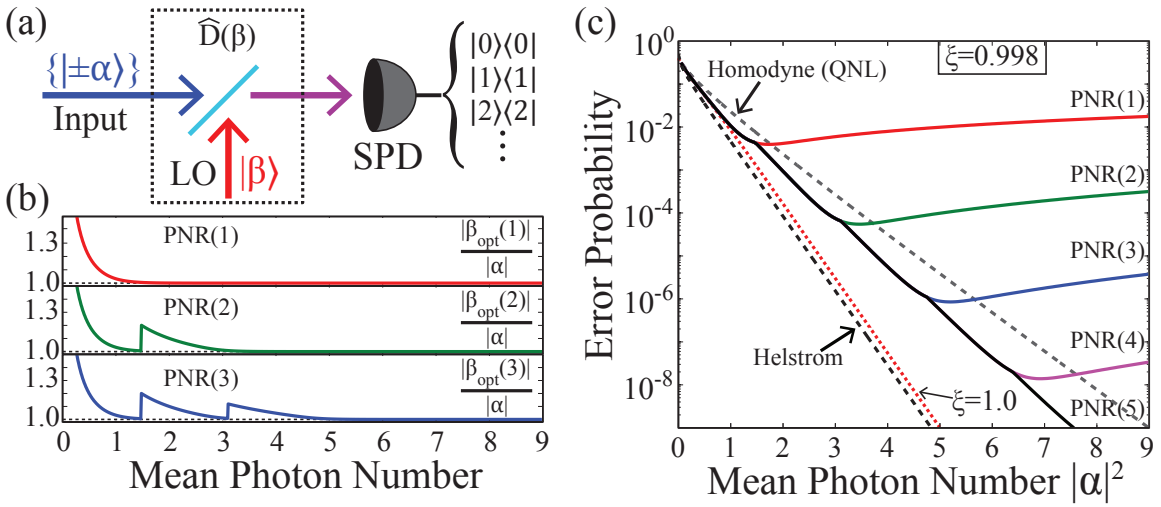

Figure 2.1: (a) The input states are displaced in phase space using interference on a beam splitter. The displaced states are then detected by a photon number resolving detector. (b) Optimal displacement ratios $|\beta_{opt}(m)|/|\alpha|$ for different PNR(m) strategies. (c) The resulting error probability (colored lines) for the receiver when implementing the optimized displacements in (b) with ideal detection efficiency, but non-ideal interference visibility ξ. The black, red, and gray dashed lines show the Helstrom bound, expected error when $\xi = 1.0$, and QNL given by an ideal homodyne measurement, respectively. Figure adapted from: *Physical Review Letters* 121, 023603 (2018).

for convenience, and to correspond with the experimental implementation, we have included an implicit π phase shift into the displacement operation. Given a particular detection outcome n, the posterior probabilities $P(\pm\alpha|n,\beta,m)$ for the input states $|\pm\alpha\rangle$ are obtained through Bayes rule: $P(\pm\alpha|n,\beta,m)P(n) = \mathcal{L}(n|\pm\alpha,\beta,m)P(\pm\alpha)$. The prior probabilities $P(\pm\alpha)$ of the input states $|\pm\alpha\rangle$ are equal to 1/2, and

$$\mathcal{L}(n|\alpha_k,\beta,m) = \frac{(\langle\hat{n}\rangle_k)^n}{n!}e^{-\langle\hat{n}\rangle_k} \tag{2.1}$$

is the likelihood function for detecting n photons for each input state given the displacement β with PNR(m). We include experimental imperfections into $\mathcal{L}(n|\alpha_k,\beta,m)$ by modeling the experimentally detected mean photon number $\langle\hat{n}\rangle_k$ of the distribution as:

$$\langle\hat{n}\rangle_k = \eta|\alpha_k-\beta|^2+\nu = \eta\big(|\alpha_k|^2+|\beta|^2-2\xi|\alpha_k||\beta|\cos(\phi_k)\big)+\nu. \tag{2.2}$$

Here η is the overall detection efficiency, ξ is the visibility and characterizes imperfections in the interference of the two fields, ν is the dark count rate of the

detector, and ϕ_k is the relative phase between the input $|\alpha_k\rangle$ and LO field (either $\phi_k = 0$ or $\phi_k = \pi$ for BPSK states) [80]. The strategy follows a maximum *a posteriori* probability (MAP) criterion such that the estimated state is the one with the largest posterior probability. Thus, there is a decision boundary that splits the $m + 1$ possible outcomes (photon detections) into two groups corresponding to which input state is more likely. For example, a PNR(4) strategy has five outcomes $\{0, 1, 2, 3, 4+\}$, which for $|\alpha|^2 = 5$ are divided as: $\{0, 1\} \rightarrow |-\alpha\rangle$, $\{2, 3, 4+\} \rightarrow |\alpha\rangle$ when the displacement $\hat{D}(\beta)$ is set to displace $|-\alpha\rangle$ to near the vacuum state. The total probability of error $P_E(\alpha, \beta, m)$ for a given displacement β and PNR(m) is [80]:

$$P_E(\alpha, \beta, m) = 1 - \frac{1}{2}\sum_{n=0}^{m} \max\big(\{P(n|\alpha, \beta, m), P(n|-\alpha, \beta, m)\}\big). \tag{2.3}$$

Given the expression for the probability of error, we find the optimal displacement magnitude $|\beta_{opt}(m)|$ by minimizing Eq. (2.3) with respect to $|\beta|$ for a fixed mean photon number $|\alpha|^2$ and resolution m. Note that by construction due to the MAP rule, the optimal decision boundary is also obtained from the minimization. Figure 2.1(b) shows the resulting optimal displacement ratio $|\beta_{opt}(m)|/|\alpha|$ for m=1, 2, and 3 with ideal detection efficiency ($\eta = 1$) but a non-ideal interference visibility of $\xi = 0.998$. The optimal displacement ratio in general converges to a value of $|\beta_{opt}(m)|/|\alpha| = 1$ (corresponding to displacement to the vacuum state) as the mean photon number increases. The jumps for $m \neq 1$ are from optimization of Eq. (2.3) and correspond to where the decision boundary changes.

Figure 2.1(c) shows the error probabilities (colored lines) for strategies with different PNR(m) for $\eta = 1$ and $\xi = 0.998$ compared to the homodyne limit (gray dashed), the Helstrom bound (black dashed), and an ideal strategy with $\xi = 1.0$ (red dashed). For an realistic experiment where $\xi = 0.998$, a strategy with PNR(1) encounters a noise floor at approximately $P_E \approx 5.0 \times 10^{-3}$ due to erroneous photon detections from the state which was displaced to *almost* the vacuum state. Increasing the PNR can counteract this effect and allow the receiver to maintain sub-QNL performance in the presence of experimental imperfections. As the input mean photon number increases, the PNR of the receiver can be further increased to provide robustness. The points

where different PNR(m) curves start to diverge correspond to the jumps in (b) of the optimal displacement ratios and changes in the decision boundary.

2.1.2 Optimization and Implementation

The first step to implementing the optimized Kennedy receiver with PNR is to optimize the value of the displacement $|\beta|$. There will in general be m local minima of Eq. (2.3) for a PNR(m) strategy, one of which being the global minimum $|\beta_{opt}(m)|$. However, we find that the particular local minima that corresponds to the global minimum changes depending on the value of $|\alpha|^2$. Figure 2.2 shows two examples of the value of $P_E(\alpha, \beta, m)$ from Eq. (2.3) as a function of the mean photon number of the displacement $|\beta|^2$. Figure 2.2(a) shows $P_E(\alpha, \beta, m)$ for $|\alpha|^2 = 1.4$ and 1.6 as dashed and solid lines for m=1, 2, and 3, in red, green, and blue, respectively. By examining the dashed lines for $|\alpha|^2 = 1.4$, we see that there are local minima at $|\beta|^2 \approx 1.5$, 2.1, and 3.0 with the global minimum (GM) for all PNR being at $|\beta|^2 \approx 1.4$. However, as the input power is increased to $|\alpha|^2 = 1.6$ the local minima shift slightly and now the

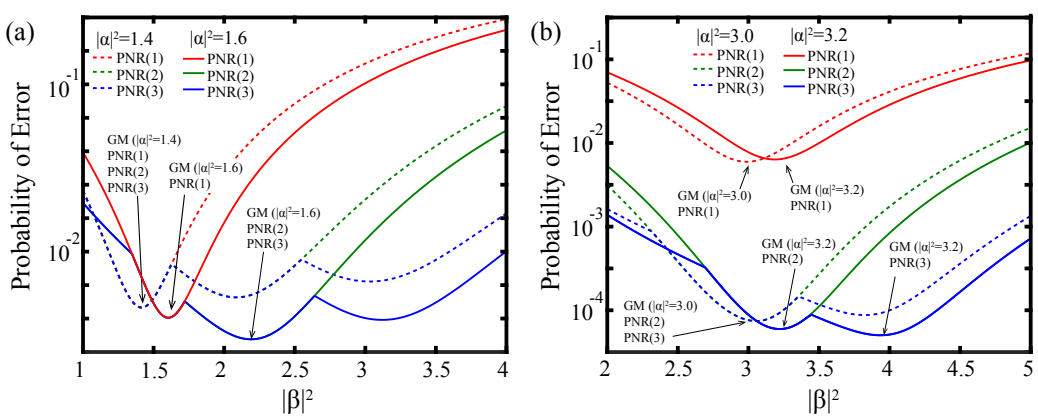

Figure 2.2: (a) Probability of error given by Eq. (2.3) for $|\alpha|^2$=1.4 (dashed) and 1.6 (solid) for m=1, 2, 3 in red, green, and blue, respectively. The global minimum (GM) for PNR(1,2,3) at $|\alpha|^2 = 1.4$ is at $|\beta|^2 \approx 1.4$ while for $|\alpha|^2 = 1.6$, the GM for PNR(1) is $|\beta|^2 \approx 1.6$ but at $|\beta|^2 \approx 2.2$ for PNR(2,3). (b) Similarly, the GM changes for PNR(2) and PNR(3) when the input energy changes from $|\alpha|^2 = 3.0$ to 3.2. Figure adapted from: *Physical Review Letters* 121, 023603 (2018).

global minimum for PNR(1) is $|\beta|^2 \approx 1.6$ but the minimum for PNR(2,3) is $|\beta|^2 \approx 2.2$. Figure 2.2(b) shows similar behavior at higher input powers of $|\alpha|^2 = 3.0$ and 3.2 where the global minimum for PNR(1) is very far from the others and the optimal value splits for PNR(2) and PNR(3). Thus, a simple function minimization may not find the global minimum depending on the initial search value. To reliably find the global minimum, we implement multiple optimizations each with different initial values. The initial values for a strategy with $|\alpha|^2$ can be roughly found empirically, or by extrapolating based on the exact local minima found for $|\alpha|^2 - \epsilon$ where ϵ is a quantity small enough such that each local minima can be "tracked" as the input energy increases.

Figure 2.3 shows the experimental setup for the optimized Kennedy receiver with PNR (see Appendix A.2 for details on how PNR is experimentally implemented). Once the optimal value of $|\beta|$ is found, the input states are prepared by adjusting the energy $|\alpha|^2$ using a series of attenuators and setting the phase randomly to either 0 or π to give the two states: $|\alpha\rangle$ and $|-\alpha\rangle$. The LO field is prepared by fixing the

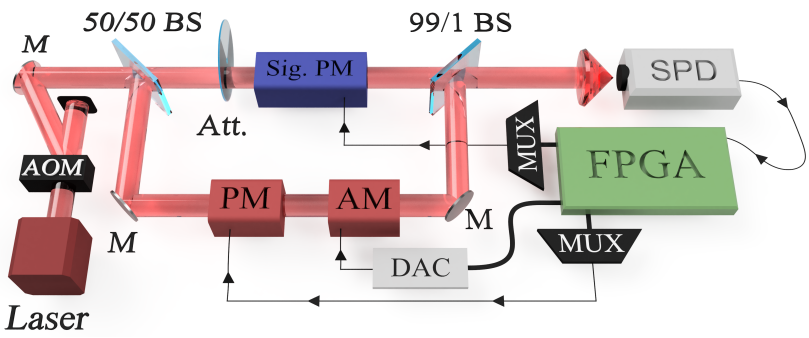

Figure 2.3: Experimental interferometer setup. Pulses of light are generated using a continuous Helium-Neon laser and a pulsed acousto-optic modulator (AOM). The pulses are split by a 50/50 beam splitter (BS) where the signal energy is prepared with an attenuator (Att.) and phase with a phase modulator (PM) and multiplexer (MUX). The LO (displacement) field is prepared with a second PM and an amplitude modulator (AM) which is connected to a digital-to-analog converter (DAC) to implement arbitrary displacement magnitudes. The two fields interfere on a 99/1 BS where the displaced signal is then detected by a single photon detector (SPD). The experiment is controlled by a field-programmable gate array (FPGA) which registers the photon detection, sets the displacement, and passes information to a computer for processing.

phase to zero and adjusting the intensity $|\beta|^2$ using an amplitude modulator and an eight bit digital-to-analog converter (DAC) [80]. Our implementation achieves an overall detection efficiency of $\eta = 0.72$, visibility of $\xi = 0.998$, dark count rate of $\nu = 3.6 \times 10^{-3}$, and APD after-pulsing probability (See Appendix B.2 for details) of $P_{AP} = 1.1 \times 10^{-2}$.

2.1.3 Experimental Results

Figure 2.4 shows the experimental results for the optimized Kennedy receiver with PNR(m) for m=1, 2, 3, and 4. The colored points and error bars correspond to the average and one standard deviation across five experimental runs and the colored dashed lines show the theoretical results obtained from Monte Carlo simulations including experimental imperfections. A comparison to a homodyne measurement with the same efficiency (dashed black) shows the robustness to noise provided by PNR as the mean photon number increases. We find that a PNR(1) strategy can only outperform the adjusted QNL up to an input energy of $|\alpha|^2 \approx 2.0$, but by increasing to PNR(2) the performance can be maintained up to $|\alpha|^2 \approx 4.5$, where an increase to PNR(3) then enables performance below the QNL, and so on.

This result shows that the optimized Kennedy receiver can always outperform a homodyne receiver for any input mean photon number provided there is a large enough photon number resolution. We note that the error probability for a given PNR(m) strategy slightly increases once it degrades relative to larger PNR strategies, which is apparent from the PNR(1) result. This effect is due to the receiver no longer being able to compensate for imperfections such as non-ideal visibility in the experiment. Thus, with PNR(1) the receiver is attempting to displace effectively to vacuum but can only reach $\hat{D}(\alpha)|\alpha\rangle = |\sqrt{2(1-\xi)}\alpha\rangle$ ($\approx |0.14\rangle$ for $|\alpha|^2 = 5$). The probability of detecting a photon from this state (and thus an error) is given by $1 - e^{-2(1-\xi)|\alpha|^2}$ (≈ 0.02 for $|\alpha|^2 = 5$) and increases as the input energy increases. One may surmise however, that after a certain point it would actually be more beneficial to encode

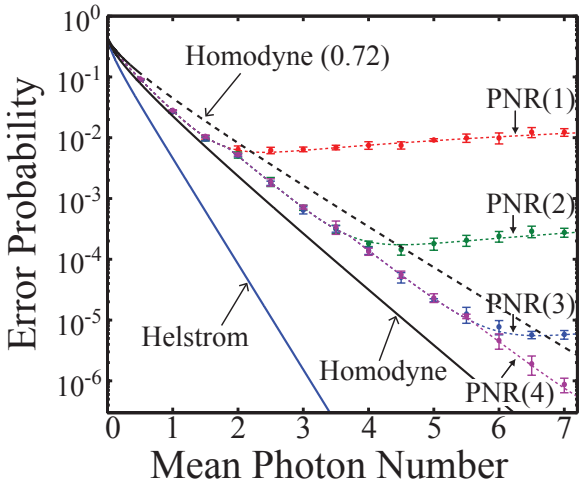

Figure 2.4: Experimental results. Points show the results for different PNR strategies and error bars correspond to one standard deviation of five separate runs of the experiment. Increasing the PNR allows the receiver to maintain its benefit over an equivalent homodyne receiver (dashed black) as the input mean photon number increases. Figure adapted from: *Physical Review Letters* 121, 023603 (2018).

information into an OOK alphabet (whose error doesn't depend on ξ) as opposed to BPSK due to this increase in error because of imperfections. Tailoring the alphabet to remain highly sensitive while providing robustness to communication channel noise is discussed in Sec. 2.2.

We further investigate the capabilities of the PNR strategies for different experimental conditions which are realistically achievable for different values of ξ and m. Figure 2.5(a-f) shows for ratio $\log_{10}(P_E/P_{hom})$ of the expected error probability to that of an ideal homodyne receiver (QNL) for PNR(1,3,5) as a function of the interference visibility ξ and mean photon number $|\alpha|^2$ for $\eta = 0.85$ in (a-c) and $\eta = 1.0$ in (d-f). The blue regions correspond to when the optimized Kennedy receiver with PNR achieves lower error rates than the QNL. The dashed green vertical line shows our experimental visibility of $\xi = 0.998$. The solid gray lines show the boundary where the error probability is no longer less than the QNL and the gray dashed line in (b,c,e,f) show this boundary for the smaller values of m. The value of $\eta = 0.85$ corresponds to using a state-of-the-art low dark count detector with 98%

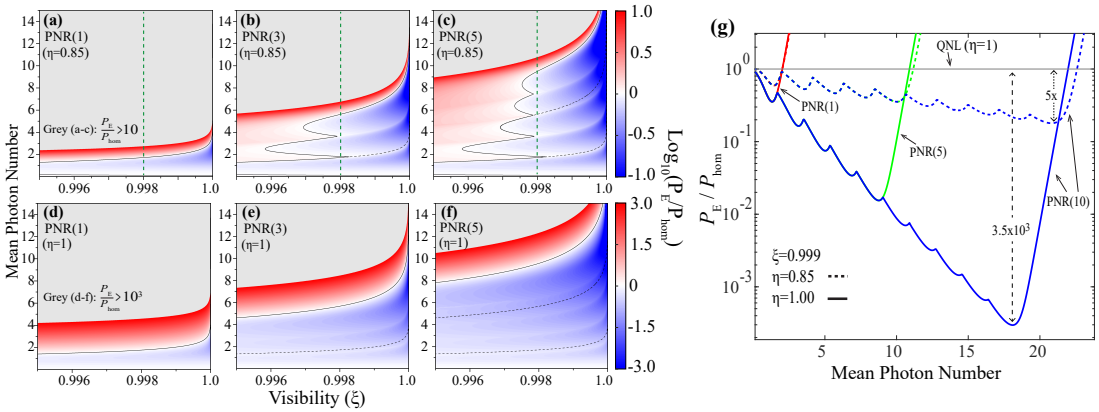

Figure 2.5: (a-f) Ratio of the expected error probability compared to the QNL $\log_{10}(P_E/P_{hom})$ as a function of visibility ξ and $|\alpha|^2$ for PNR(1,3,5) with $\eta = 0.85$ in (a-c) and $\eta = 1.0$ in (d-f). The dashed green vertical line shows our experimental visibility of $\xi = 0.998$ and the solid gray lines show where $P_E = P_{hom}$. (g) Comparison of different strategies to the ideal QNL for an experimentally feasible visibility of $\xi = 0.999$ for $\eta = 0.85$ (dashed) and $\eta = 1.0$ (solid). A PNR(10) strategy with $\eta = 0.85$ can achieve 5 times better performance than the ideal QNL and a perfect $\eta = 1.0$ strategy can reach over 10^3 times better than the QNL. Figure adapted from: *Physical Review Letters* 121, 023603 (2018).

efficiency combined with our total system efficiency of $\eta_{system} \approx 0.88$. In this practical case with a realistic visibility, the receiver can maintain its performance relative to the QNL with PNR(5) up to $|\alpha|^2 \approx 10$. Interestingly, sub-QNL performance can be obtained for $|\alpha|^2 < 2$ even as the visibility decrease beyond $\xi = 0.995$. Figure 2.5(g) shows the ratio of the two error probabilities for an experimentally feasible visibility of $\xi = 0.999$ for $\eta = 0.85, 1.0$ and PNR$(1, 5, 10)$. A PNR(10) strategy which a realistic photon number resolving detector with $\eta = 0.85$ can achieve approximately five times better performance relative to the ideal QNL and the same strategy with a perfect detector can reach error rates of over 10^3 times lower than the QNL.

2.1.4 Discussion

The demonstration of optimized strategies with PNR solves the important issue of robustness in Kennedy-like receivers for discrimination of BPSK coherent states. Whereas previous demonstrations of binary state discrimination below the QNL have

been limited to low input powers due to noise an imperfections, we show that these effects can be mitigated by using an adequate photon number resolution. Increasing the PNR from simple on/off detection provides robustness to non-ideal interference visibility as well as detector dark counts, and after-pulsing effects. Allowing the optimized Kennedy receiver to have photon number resolution can also be leveraged for quantum key distribution with coherent states [98], by allowing the receiver to implement a more "soft" decision strategy as opposed to a "hard" decision with on/off detection.

2.2 Optimized Communication over Phase Noise Channels

In addition to noise and imperfections in the measurement receiver, there can potentially be noise imparted onto the coherent states by the communication channel itself. In such a scenario, the attainable performance of the state discrimination measurement will generally be degraded, especially at high levels of communication channel noise. For certain types of channel noise such as loss or phase insensitive Gaussian noise, coherent states of light can reach the limits of information transfer, known as the capacity [18, 71, 99]. However, the ultimate limits of communication [72–75, 95, 100–102] and metrology [103–105] in non-Gaussian channels, and strategies which achieve those limits, are still under investigation. To understand the impact of non-Gaussian channels on state discrimination, experimental investigations of strategies which perform well over these channels are needed. To this end, we study and experimentally demonstrate an optimized communication strategy for binary coherent states over a phase diffusion channel utilizing the optimized Kennedy receiver with PNR described above [81]. The communication strategy jointly optimizes the measurement receiver as well as the input alphabet to the channel under an average energy constraint. We show that this approach can achieve lower error probabilities compared to conventional approaches, as well as reach higher rates of information

transfer for small and large noise levels.

2.2.1 Optimized Communication Strategy

A central aspect of the optimized communication strategy is the ability to tailor the binary input state alphabet to reduce the impact of the phase noise channel on the states. Figure 2.6 shows the effect of a phase diffusion channel on three different input coherent state alphabets with the same *average* energy $\bar{n}$: 2.6(a) shows a BPSK alphabet $\{|-\alpha\rangle, |+\alpha\rangle\}$, 2.6(b) shows an OOK alphabet $\{|0\rangle, |\sqrt{2}\alpha\rangle\}$, and 2.6(c) shows a general alphabet $\{|\alpha_1\rangle, |\alpha_2\rangle\}$ where $|\alpha_1|^2 + |\alpha_2|^2 = 2\bar{n}$. A BPSK alphabet achieves the smallest possible overlap $|\langle -\alpha|\alpha\rangle|^2$ for a given average energy $\bar{n}$, but it

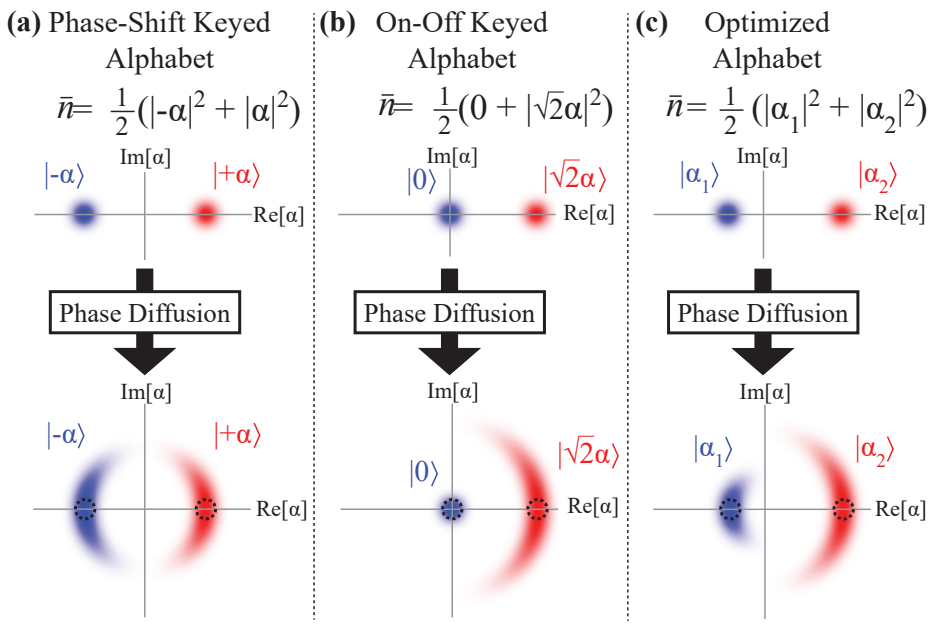

Figure 2.6: Effect of phase diffusion on (a) a BPSK alphabet, (b) an OOK alphabet, and (c) an optimized alphabet. The BPSK alphabet achieves the best error rate in the absence of phase diffusion, but it is it most sensitive as the overlap between the states quickly increases with phase diffusion. On the other hand, an OOK alphabet has higher error rates with no phase diffusion but is also immune to this type of channel noise. An optimized alphabet seeks to balance the benefits of BPSK and OOK to achieve the best possible error in the presence of a particular level of phase diffusion. Figure adapted from: *npj Quantum Information* 5, 65 (2019).

is highly sensitive to phase diffusion, which decreases the distinguishability of the states [2]. On the other hand, an OOK alphabet has a larger overlap when there is no noise and therefore a larger error probability. However, OOK is immune to the effects of phase diffusion as information is encoded into the magnitude of the states rather than the phase. An optimized alphabet captures both the high distinguishability of BPSK and robustness to phase diffusion of OOK due to the ability to interpolate between the two alphabets for a given level of channel noise.

This idea of an optimal alphabet is combined with the previously discussed optimized Kennedy receiver with PNR such that both the alphabet and receiver are *jointly* optimized. Figure 2.7(a) shows the optimized communication strategy over a phase diffusion channel. The sender prepares a state from the alphabet $\{|\alpha_k\rangle\}$ for $k = 1, 2$ and sends the state through the phase diffusion channel. The effect of the channel noise transforms the input alphabet into:

$$\hat{\rho}_k = |\alpha_k\rangle\langle\alpha_k| \to \hat{\rho}_k(\sigma) = \int_{-\infty}^{\infty} \frac{e^{-\frac{\phi^2}{2\sigma^2}}}{\sqrt{2\pi\sigma^2}} |\alpha_k e^{-i\phi}\rangle\langle\alpha_k e^{-i\phi}| d\phi, \tag{2.4}$$

where the phase diffusion noise strength is determined by σ. The receiver implements an optimized Kennedy receiver with PNR to discriminate the phase diffused states. The alphabet and displacement operation $\hat{D}(\beta)$ are jointly optimized such that the error probability is minimized for a given value of σ, average mean photon number $\bar{n}$, and PNR(m). Similar to Eq. (2.3), the error probability for equiprobable states where $P(\hat{\rho}_{1,2}(\sigma)) = 0.5$ can be constructed using the MAP criteria as:

$$P_E(\bar{n}, \{\hat{\rho}_{1,2}(\sigma)\}, \beta, m) = 1 - \frac{1}{2} \sum_{n=0}^{m} \max_k \left(\{\mathcal{L}(n|\hat{\rho}_k(\sigma), \beta, m)\}\right). \tag{2.5}$$

This error probability depends on the displacement β, resolution m, the energy distribution of $|\alpha_1|^2$ and $|\alpha_2|^2$, and phase diffusion strength σ through Eq. (2.4).

Figure 2.7(b) shows error probability of the optimized communication strategy (solid blue) for $\bar{n} = 0.5$ compared to a BPSK strategy (dashed red) as the the level of phase noise increases. The black line shows the Helstrom bound $P_E^{Hel} =$

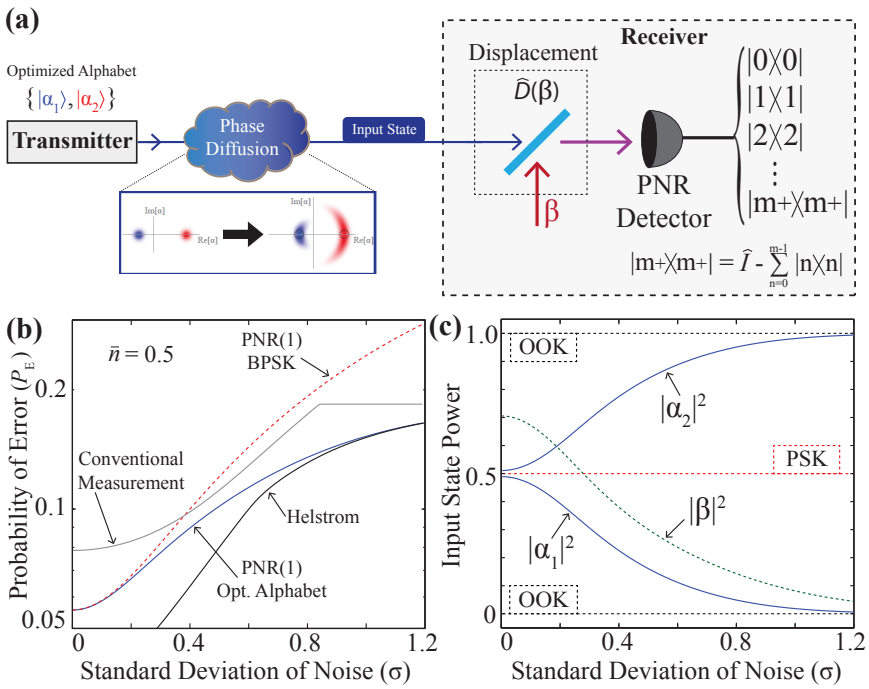

Figure 2.7: (a) Optimized communication strategy where the transmitter prepares an optimal alphabet which undergoes phase diffusion due to a non-Gaussian channel. The receiver implements an optimized Kennedy measurement with PNR(m) which is optimized jointly with the input alphabet to minimize the overall error probability. (b) Error rate as a function of the noise level σ for an optimized strategy (blue), a BPSK encoding (red dashed), Helstrom bound (black) with its own optimized alphabet, and a conventional measurement (gray) also with its own optimal alphabet. (c) Energy of the optimal input states ($|\alpha_1|^2, |\alpha_2|^2$) and displacement level ($|\beta|^2$). Figure adapted from: *npj Quantum Information* 5, 65 (2019).

$\frac{1}{2}\left(1 - \frac{1}{2}||\hat{\rho}_1(\sigma) - \hat{\rho}_{1,2}(\sigma)||_1\right)$ with the alphabet $\{\hat{\rho}_1(\sigma), \hat{\rho}_2(\sigma)\}$ which minimizes this bound and in general will be different from the alphabet which minimizes the error for a particular measurement. The gray line shows the error probability of what we refer to as a conventional measurement (CM), which is either a homodyne measurement with its own optimized alphabet or direct detection with OOK, whichever has the lower error. We find that the main benefit of the optimized communication strategy is that by modifying the input alphabet, the strategy can interpolate between a BPSK alphabet when there is little noise and an OOK alphabet for higher noise levels. An interesting result of this interpolation is that the optimized strategy asymptotically approaches the Helstrom bound as the phase diffusion increases beyond $\sigma \approx 1$.

Figure 2.7(c) shows this optimal alphabet (solid blue) for the strategy as a function of the phase noise level, where the interpolation between BPSK and OOK can be explicitly seen, as well as the optimal displacement energy $|\beta|^2$ (dashed green). We note that even when $\sigma = 0$, the optimal alphabet is not *exactly* a BPSK alphabet due to the incorporation of a non-ideal interference visibility of $\xi = 0.998$, which the strategy can compensate for by slightly adjusting the alphabet even in the absence of channel noise. This would imply that optimizing the communication alphabet can provide robustness to experimental imperfections, such as visibility and dark counts, as well as channel noise.

2.2.2 Experimental Results

Error Probability for BPSK and Optimized Alphabets

We experimentally demonstrate the optimized communication strategies described above using the experimental setup described in Sec. 2.1.2. We implement the phase diffusion by applying a voltage waveform which is piecewise constant and Gaussian distributed to the signal phase modulator. We calibrate the amplitude of the voltage waveform following the procedure described in Ref. [81], where each amplitude level for the waveform is associated with a corresponding standard deviation σ of the measured phase shifts from the waveform. In order to first benchmark the performance of the optimized communication scheme, we implement a BPSK alphabet with an optimized Kennedy receiver with PNR. Figure 2.8(a-c) shows the experimental results for the BPSK alphabet for an average input mean photon number of $\bar{n} = 0.5$, 1.0, and 2.0, respectively. For $\bar{n}$=0.5, we find that the receiver can tolerate a noise level of $\sigma \approx 0.6$ before the error increases above an equivalent homodyne measurement. For $\bar{n}$=2.0, the benefit of PNR for robustness to channel noise becomes apparent and while PNR(1) can only tolerate noise up to $\sigma \approx 0.06$, increasing to PNR(3) allows the receiver to outperform a homodyne measurement up to $\sigma \approx 0.25$. However for average energies, as the level of noise increases the achievable error probability asymptotically approaches

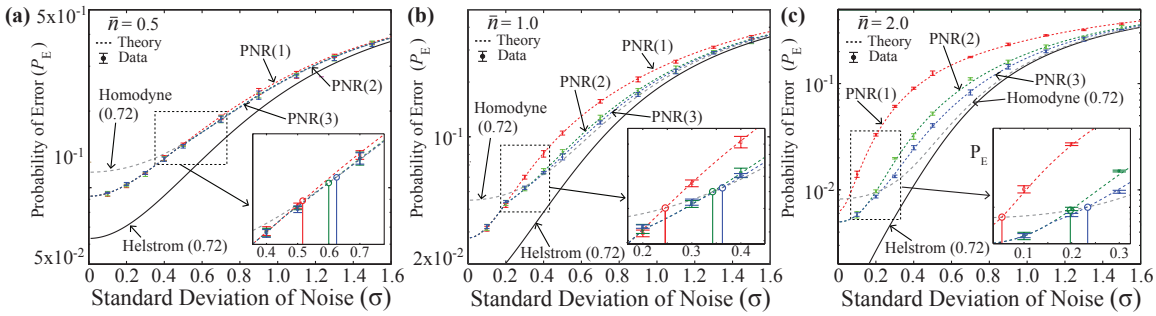

Figure 2.8: Experimental error probabilities for an optimized Kennedy receiver with a BPSK alphabet for $\bar{n}$=0.5, 1.0, and 2.0 in (a-c), respectively. For $\bar{n} = 0.5$ in (a), the receiver can only maintain an error probability better than an equivalent homodyne receiver up to a noise level of $\sigma \approx 0.6$ and the use of photon number resolution has a small impact. For $\bar{n} = 2.0$, the effect of increasing from PNR(1) to PNR(3) becomes apparent and can increase the tolerable noise level by a factor of ≈ 5. Figure adapted from: *npj Quantum Information* 5, 65 (2019).

0.5 as the overlap of the states approaches unity.

Figure 2.9(a-c) shows the error probability for the optimized communication strategy for $\bar{n}$=0.5, 1.0, and 2.0, respectively, and 2.9(d-f) shows the corresponding optimal input alphabet. For $\bar{n} = 0.5$, we find that the receiver is able to outperform a conventional measurement (dashed gray) strategy for all levels of phase noise σ and even approach the equivalent Helstrom bound (solid black) as σ increases. The red and green points correspond to PNR(1) and PNR(3) strategies and the error bars represent one standard deviation across five experimental runs. For $\bar{n} = 2.0$, we see the combined effect of an optimized communication strategy and PNR on the probability of error. While a PNR(1) optimized alphabet can outperform BPSK, the error is significantly above a conventional measurement for a moderate amount of noise. However, the optimized alphabet with PNR(3) can leverage the number resolving capabilities such that there is only a small range of phase noise for which the error of the optimized strategy is slightly higher than that of a conventional measurement. The optimal alphabet is also able to remain closer to BPSK than OOK with an increase in PNR as in (f) such that the sensitivity of a coherent encoding is fully utilized. Overall, the alphabet tends to OOK as σ increases, but we note that different PNR can have

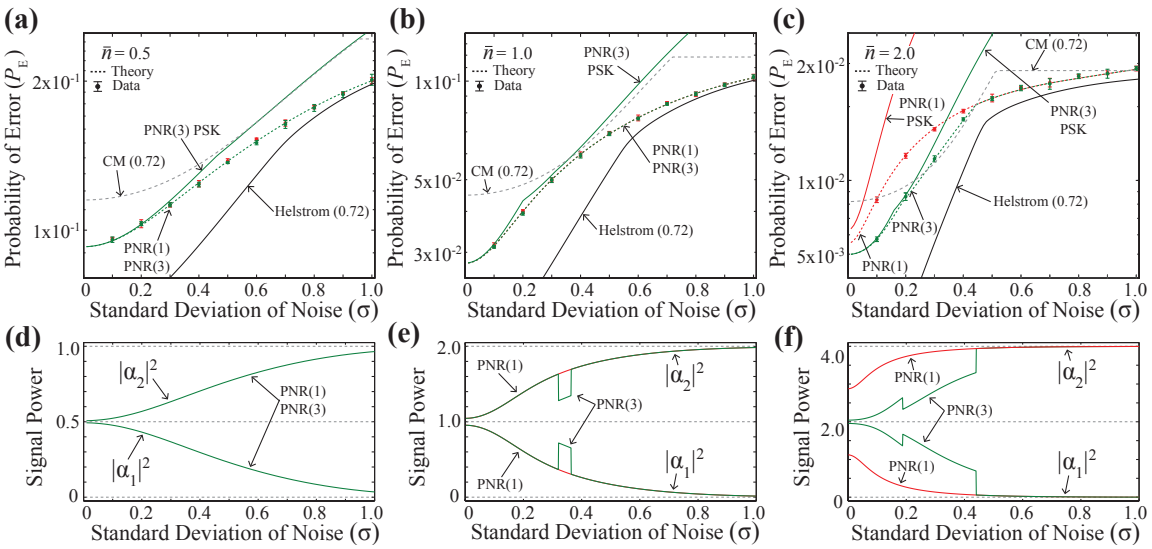

Figure 2.9: Error probabilities (a-c) and resulting alphabets (d-f) for the optimized communication strategy with $\bar{n}$=0.5, 1.0, and 2.0, respectively. In (a), the strategy can outperform an equivalent conventional measurement (CM) for all noise levels and approaches the Helstrom bound due to the ability of the alphabet to interpolate between BPSK and OOK. In (c), increasing from PNR(1) to PNR(3) allows the receiver to perform better than an equivalent CM for almost all values of σ and the corresponding optimal alphabet shows jumps similar to previous experiments comparing different number resolutions. Figure adapted from: *npj Quantum Information* 5, 65 (2019).

very different optimal alphabets and displacement values.

Interestingly, in Figure 2.9(c) for PNR(1) the error of the optimized alphabet is lower than for BPSK when there is no channel noise ($\sigma = 0$). This effect is due to the optimal alphabet attempting to compensate for imperfections in the experimental implementation, mainly a non-ideal visibility of $\xi = 0.998$. This effect is apparent by examining the optimal alphabet in (f) where $|\alpha_1|^2 \approx 1.0$ and $|\alpha_1|^2 \approx 3.0$ instead of $|\alpha_1|^2 = |\alpha_2|^2 = 2.0$ for BPSK. Thus, for realistic implementations with limited PNR, an optimized alphabet can further improve the error probability by taking experimental imperfections in to account in the optimal encoding, even in just a pure loss channel ($\sigma = 0$).

Mutual Information for Optimized Alphabets

We further investigate the effect of a phase diffusion channel on information transfer by studying the mutual information (MI) $I(X:Y)$ for optimized communication strategies with a phase diffusion channel. The mutual information quantifies the amount of information (in bits) which can be reliably transmitted across a channel using a particular encoding and measurement strategy. The maximum MI for binary states is one bit, and we aim to find communication strategies which extract the highest possible amount of information. We calculate the MI for a displaced photon counting strategy using a "soft-decision" rule based on the number of detected photons, as opposed to a binary outcome of a particular state:

$$I(\bar{n},\sigma,\beta,m)=\sum_{n=0}^{m}\sum_{k=1}^{2}\mathcal{L}(n|\hat{\rho}_k(\sigma),\beta,m)P(\hat{\rho}_k(\sigma))\log_2\left[\frac{\mathcal{L}(n|\hat{\rho}_k(\sigma),\beta,m)}{P(n|m)}\right], \quad (2.6)$$

where $\mathcal{L}(n|\hat{\rho}_k(\sigma),\beta,m)$ is the likelihood function for detecting n photons given the input state $\hat{\rho}_k(\sigma)$, displacement β, and PNR(m), and $P(n|m)=\sum_k \mathcal{L}(n|\hat{\rho}_k(\sigma),\beta,m)P(\hat{\rho}_k(\sigma))$ is the total probability of detecting n photons with PNR(m). Similar to the error probability, the MI can be maximized and will in general have different values for the optimal alphabet and displacement for a given average energy $\bar{n}$, PNR(m), and noise level σ.

Figure 2.10(a,b) show the experimental MI (points) for different PNR for $\bar{n}=1.0$ and $\bar{n}=2.0$, respectively, and 2.10(c,d) show the corresponding alphabets. Dashed colored lines show the MI for different PNR and the dashed gray line shows the MI for a conventional measurement with its own optimized alphabet. The solid red and blue lines show the MI for a BPSK alphabet for PNR(1) and PNR(5), respectively. We observe that having a PNR greater than PNR(1) allows for higher MI than a conventional measurement at small values of σ, and all strategies can outperform conventional approaches due to the optimized alphabet at large values of σ.

Increasing the PNR can allow the optimized communication strategy to approach the MI of conventional measurements for moderate values of σ. This effect is mainly

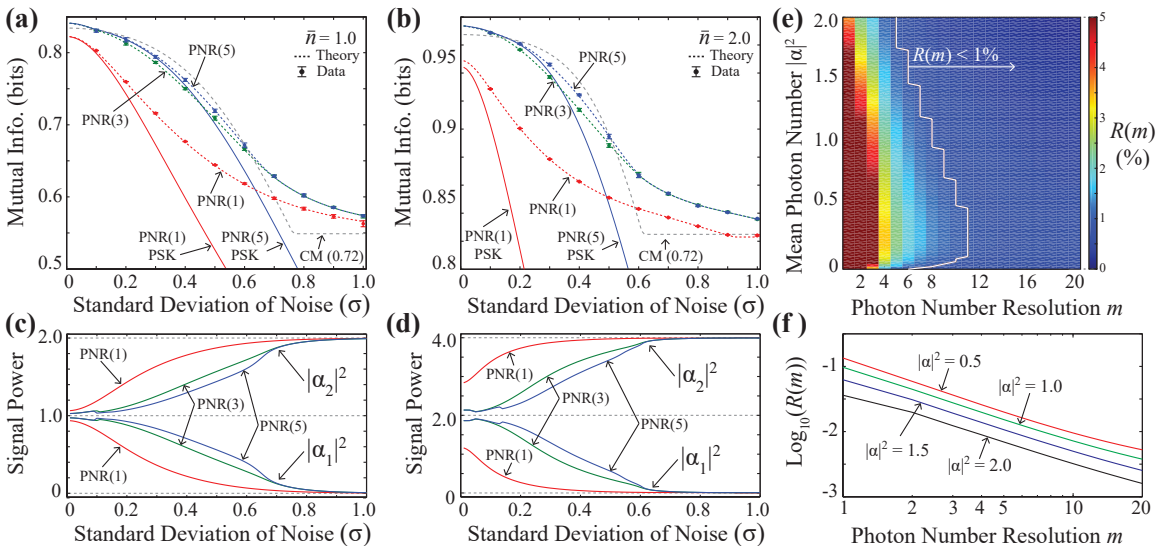

Figure 2.10: Experimentally measured mutual information (a,b) and corresponding alphabets (c,d) for $\bar{n}$=1.0 and 2.0, respectively. For both average energies, the optimized communication strategy can obtain higher mutual information in the low noise and high noise regimes, and in the intermediate regime the use of PNR greatly enhances the achievable information transfer. (e) Maximum percent difference $R(m)$ between the optimized communication strategy and an equivalent conventional measurement for different PNR and input energies. (f) $R(m)$ as a function of m for different alphabet energies showing the power-law scaling of the maximum percent difference with the PNR. Figure adapted from: *npj Quantum Information* 5, 65 (2019).

due to a PNR(1) strategy not having access to information gained from different photon number detections, rather just the presence or absence of photons. We calculate the maximum percent difference $R(m)$ between the displaced photon counting strategy and the CM for a given PNR(m):

$$R(m) = \max_{\sigma} \left(\frac{I_{CM}(\sigma) - I_{PNR}(\sigma)}{I_{CM}(\sigma)} \right), \tag{2.7}$$

where $I_{PNR}(\sigma)$ is the mutual information for the photon counting strategy and $I_{CM}(\sigma)$ is the mutual information of the CM. This value $R(m)$ represents the percent difference between the two strategies at the level of noise σ for which the a photon counting strategy has the worst performance relative to the CM. Figure 2.10(e) shows $R(m)$ for $\bar{n}$ ranging from 0 to 20 and different PNR(m) for $m = 1$ to $m = 20$. The blue region to the right of the white line corresponds to when $R(m) < 1\%$, i.e. when a PNR(m) strategy is at worst less than 1% below the CM. Figure 2.10(f) shows $R(m)$ on a

log-log scale for $\bar{n} = 0.5$, 1.0, 1.5, and 2.0, in red, green, blue, and black, respectively. The straight lines indicate power-law scaling of $R(m)$ of the form $a(m)^b$ and we find that for all $\bar{n}$, $b \approx 1.1$. This scaling indicates that for moderate phase diffusion, PNR strategies will at worst asymptotically provide the same mutual information as the CM. On the other hand, for small and large values of σ, optimized photon counting strategies with realistic PNR provide advantages in the mutual information compared to the CM.

2.2.3 Discussion

Optimized communication strategies can further our understanding of the limits of information transfer across channels with non-Gaussian noise. In particular, when these limits are not completely understood even in theory, experiments must probe how to achieve efficient and effective information transfer. Furthermore, the idea of jointly optimizing the input alphabet and receiver measurement is not restricted to binary coherent states. For example, this approach can be used to study higher order modulation formats such as QPSK or 16-QAM across phase diffusion channels with different measurement schemes. Such investigations will shed light on the underlying physics involved in non-Gaussian channels and how to approach the limits of communication across such channels.

2.3 Optimally Inconclusive Measurements

The optimized Kennedy receiver described in Sections 2.1 and 2.2 is a measurement strategy referred to as a minimum-error state discrimination (MESD) measurement. In MESD strategies, the receiver aims to minimize the probability of error P_E and the minimum achievable value is the Helstrom bound P_E^{Hel} [2, 66, 106]. Beyond the paradigm of MESD measurements, there is a different class of strategies which allow for perfect ($P_E = 0$) discrimination but at the cost of a high probability of an inconclusive

result P_I. These strategies are referred to as unambiguous state discrimination (USD) measurements where the smallest probability of an inconclusive result is given by P_I^{USD} [26, 49, 96, 97, 107, 108]. Reference [53] demonstrated a Dolinar receiver (MESD strategy) for intensity encoded coherent states, and Reference [109] demonstrated USD measurements of arbitrary binary coherent states.

While MESD and USD represent two different measurement paradigms, a more general state discrimination strategy is one which is *optimally inconclusive* [17, 110]. This more general discrimination strategy operates on the principle that for a given probability of an inconclusive result $0 \leq P_I \leq P_I^{USD}$, the receiver achieves the minimum possible error probability $0 \leq P_E \leq P_E^{Hel}$ [111, 112]. In this way, an optimally inconclusive measurement receiver has the ability to interpolate between the two canonical measurement paradigms of MESD and USD. Inconclusive measurements of coherent states have been experimentally demonstrated [98]. However, the demonstration in Ref. [98] falls short of achieving the optimality of the optimally inconclusive strategy, i.e. obtaining the lowest error for a given value of inconclusive probability. Recent theoretical work in Ref. [63] shows that optimally inconclusive measurements of binary coherent states can be implemented with a *modified* Dolinar approach using displacement operations, photon counting, and feedback. The measurement is performed by splitting the input state into two temporal modes and implementing a standard Dolinar receiver in the first temporal mode to discriminate between the states with minimal error, and then implementing a Dolinar-*like* receiver in the second temporal mode to decide between the most likely state after the first mode and an inconclusive result.

2.3.1 Measurement Strategy

The measurement operators (POVM elements) which implement the MESD and USD measurement schemes for the binary states $\{|\psi_1\rangle, |\psi_2\rangle\}$ can be represented in a two dimensional space. Figure 2.11 shows the vector representation of the two non-

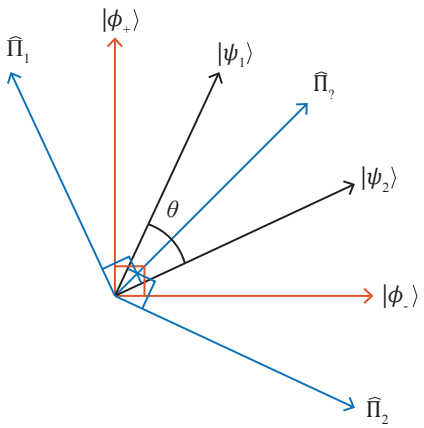

Figure 2.11: Vector space representation of the binary states and measurements. The two states $|\psi_1\rangle$ and $|\psi_2\rangle$ (black) are separated by an angle θ, which characterizes their non-orthogonality. In this two dimensional space, the MESD (Helstrom) POVM is formed by projecting onto the orthonormal vectors $|\phi_+\rangle$ and $|\phi_-\rangle$ (red), which span the two dimensional space. A USD POVM is formed by projecting onto $\hat{\Pi}_1$, $\hat{\Pi}_2$, and $\hat{\Pi}_?$.

orthogonal states $|\psi_{1,2}\rangle$ (black), the POVM for a minimum error measurement (red), and the POVM for an unambiguous measurement (blue). The optimal MESD strategy (Helstrom measurement) achieves P_E^{Hel} and is given by the POVM $\{\hat{\Pi}_1, \hat{\Pi}_2\}$, which is formed by projecting onto the orthogonal basis states $|\phi_\pm\rangle$ with $|\psi_1\rangle = a|\phi_+\rangle + b|\phi_-\rangle$ and $|\psi_2\rangle = b|\phi_+\rangle + a|\phi_-\rangle$. Here, a positive outcome of $\hat{\Pi}_1$ indicates that the more likely input state is $|\psi_1\rangle$ and vice versa. A Dolinar receiver also achieves the Helstrom bound using the experimentally implementable techniques of displacement operations, photon counting, and fast feedback. On the other hand, an unambiguous measurement is given by the POVM $\{\hat{\Pi}_1, \hat{\Pi}_2, \hat{\Pi}_?\}$ where the outcome $\hat{\Pi}_?$ represents the inconclusive outcome, which occurs with probability P_I. The USD measurement is constructed by having $\hat{\Pi}_1$ be orthogonal to $|\psi_2\rangle$ (i.e. $\langle\psi_2|\hat{\Pi}_1|\psi_2\rangle = 0$) such that a positive outcome of $\hat{\Pi}_1$ unambiguously eliminates $|\psi_2\rangle$ as a possible input state. In a similar way, $\hat{\Pi}_2$ is orthogonal to $|\psi_1\rangle$ (i.e. $\langle\psi_1|\hat{\Pi}_2|\psi_1\rangle = 0$) and a positive outcome of $\hat{\Pi}_2$ unambiguously eliminates $|\psi_1\rangle$. Thus, there are zero errors for the USD measurement ($P_E = 0$) due to how $\hat{\Pi}_1$ and $\hat{\Pi}_2$ are constructed. The measurement operator for the inconclusive outcome is given by $\hat{\Pi}_? = \hat{I} - \hat{\Pi}_1 - \hat{\Pi}_2$ and the corresponding inconclusive probability is $P_I = 1 - P_C$ where $P_C = p_1\langle\psi_1|\hat{\Pi}_1|\psi_1\rangle + p_2\langle\psi_2|\hat{\Pi}_2|\psi_2\rangle$. An optimal USD measurement

of binary coherent states $\{|\psi_{1,2}\} = \{|\alpha_{1,2}\rangle\}$ can be implemented using displacement operations and photon counting without feedback, and obtains the smallest possible inconclusive probability: $P_I^{USD} = 1 - |\langle\alpha_1|\alpha_2\rangle|$ [96].

An extension to the MESD and USD paradigms is an optimally inconclusive measurement, which aims to achieve the smallest possible error probability $P_E \leq P_E^{Hel}$ for a specified inconclusive probability $P_I \leq P_I^{USD}$, and interpolates between a MESD and USD measurement. The two measurement operators $\hat{\Pi}_1$, $\hat{\Pi}_2$ for an optimally inconclusive measurement are no longer constructed to be orthogonal to each input state, resulting in $P_E \neq 0$. Thus, the measurement probabilities P_C, P_E, and P_I are given by:

$$P_C = p_1\langle\psi_1|\hat{\Pi}_1|\psi_1\rangle + p_2\langle\psi_2|\hat{\Pi}_2|\psi_2\rangle, \tag{2.8}$$

$$P_E = p_1\langle\psi_1|\hat{\Pi}_2|\psi_1\rangle + p_2\langle\psi_2|\hat{\Pi}_1|\psi_2\rangle, \tag{2.9}$$

$$P_I = p_1\langle\psi_1|\hat{\Pi}_?|\psi_1\rangle + p_2\langle\psi_2|\hat{\Pi}_?|\psi_2\rangle = 1 - P_E - P_C. \tag{2.10}$$

Sub-optimal inconclusive measurements of binary coherent states have been experimentally demonstrated [98], but do not achieve minimal P_E given the inconclusive probability P_I. Furthermore, recent theoretical work [63] shows that an optimally inconclusive measurement for coherent states with arbitrary prior probabilities can be implemented by using a *modified* Dolinar receiver.

Figure 2.12(a) shows the Helstrom bound on the error probability P_E (orange) for equiprobable BPSK coherent states ($|\psi_1\rangle = |\alpha\rangle$, $|\psi_2\rangle = |-\alpha\rangle$), and the conclusive probability ($1-P_I^{USD}$) of an optimal unambiguous measurement (blue) as a function of the input energy of the states $|\alpha|^2$. Figure 2.12(b) shows the theoretical measurement probabilities (P_I, P_E) of the optimally inconclusive receiver for equiprobable BPSK coherent states with input mean photon numbers of $|\alpha|^2 = 0.2$, 0.4, and 0.6, in blue, orange, and yellow, respectively. This result is also obtained by relating the non-orthogonality of the two coherent states ($|\langle-\alpha|\alpha\rangle|^2$) to an angle θ between $|\psi_1\rangle$ and $|\psi_2\rangle$ in Fig. 2.11 with the equation: $e^{-4|\alpha|^2} = \cos^2(\theta)$.

The standard Dolinar receiver for BPSK coherent states is implemented using

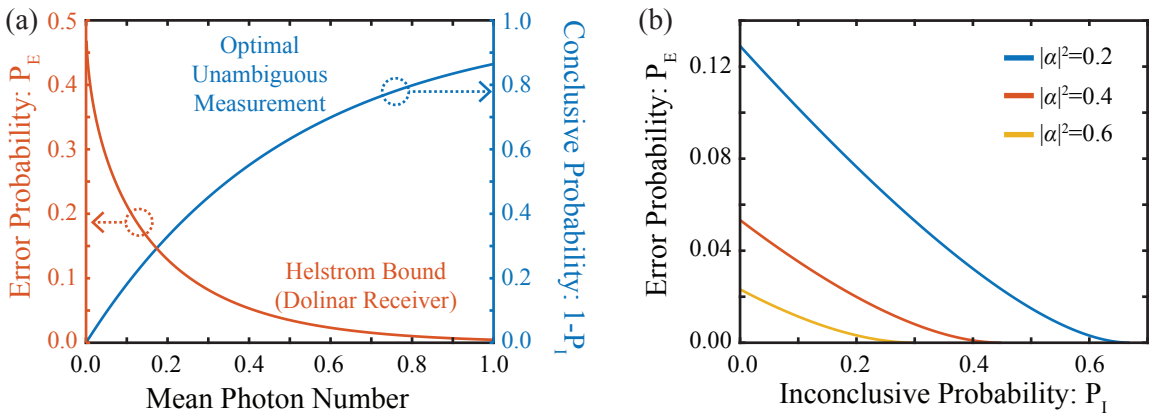

Figure 2.12: (a) Error probability for the Dolinar receiver performing at the Helstrom bound for MESD and the conclusive probability $(1 - P_I^{USD})$ from a receiver performing a USD measurement. (b) Theoretical results for the error probability P_E as a function of the inconclusive probability P_I for $|\alpha|^2$=0.2, 0.4, and 0.6 for an optimally inconclusive measurement.

displacement operations, photon counting, and feedback. The input state is displaced in phase space using interference with a LO field, where the magnitude of the LO has an optimal waveform $|u(t)|$ and the phase of the LO (sign of $u(t)$) is switched between 0 and π every time a photon is detected [53,62]. The authors in Ref. [63] show that an optimally inconclusive measurement for coherent states can be implemented with a *modified* Dolinar scheme. In this modified scheme, the input state is split into two temporal modes and measured using displacement operations, photon counting, and fast feedback, with a particular waveform for the displacement operation. The duration of the input state is normalized such that $0 \leq t \leq 1$, where the first temporal mode has a duration of $0 \leq t \leq t_1$ and the second temporal mode has a duration of $t_1 < t \leq 1$. The optimally inconclusive measurement from Ref. [63] implements a standard Dolinar receiver in the first temporal mode to discriminate between $|\pm\alpha\rangle$ with minimal error using the displacement waveform:

$$u_1(t) = \frac{(-1)^{N_1(t)}\alpha}{\sqrt{1 - 4p(1-p)K^{2t}}}, \tag{2.11}$$

where without loss of generality, p is the prior probability of the more likely state. The quantity $K^2 = |\langle -\alpha|\alpha\rangle|^2 = e^{-4|\alpha|^2}$ is the overlap of the two states. The total number of detected photons $N_1(t)$ $(N_1(0) = 0)$ in the first mode up to time t switches

the sign of the displacement (phase of the LO) every time a photon is detected. The instantaneous hypothesis for the input state at time t can be obtained as: $|\alpha\rangle$ if $N_1(t)$ is even and $|-\alpha\rangle$ if $N_1(t)$ is odd. This optimal waveform $u_1(t)$ can be obtained based on a Bayesian updating approach [113, 114], or from an optimal control approach [90]. From a Bayesian viewpoint, the evolution of the probability of correct detection $P_C(t)$ during the first temporal mode can be written to first order in δt as:

$$P_C(t+\delta t) = P_C(t)\Big(1 - |u_1(t) - \alpha|^2 \delta t\Big) + \Big(1 - P_C(t)\Big)|u_1(t) + \alpha|^2 \delta t, \tag{2.12}$$

where $P_E(t) = 1 - P_C(t)$ and $P_C(0) = p$. Thus, the optimal waveform $u_1(t)$ is one which maximizes $P_C(t)$ at every infinitesimal time step δt.

The optimally inconclusive receiver implements the measurement of the first temporal mode during the time interval $0 \leq t \leq t_1$. The "switching time" t_1 is obtained from the inconclusive probability P_I, prior probability p, and $|\alpha|^2$ as:

$$t_1 = \frac{1}{4|\alpha|^2} \ln\left(\frac{2p(1-p)}{2C^2 - 2CP_I + P_I}\right), \tag{2.13}$$

where $C^2 = \epsilon K^2$ and $\epsilon = p(1-p)$. Without loss of generality, the probabilities for two input states after the first temporal mode ($0 \leq t \leq t_1$) are $\{P_C^{(1)}, 1 - P_C^{(1)}\}$ where:

$$P_C^{(1)} = \frac{1}{2}\Big(1 + \sqrt{1 - 4p(1-p)e^{-4|\alpha|^2 t_1}}\Big). \tag{2.14}$$

Thus, $P_C^{(1)}$ corresponds to the Helstrom bound for the input states $\{|\sqrt{t_1}\alpha\rangle, |-\sqrt{t_1}\alpha\rangle\}$.

The prior probabilities for the input states at the beginning the second temporal mode ($t_1 < t \leq 1$) are then $\{P_C^{(1)}, 1 - P_C^{(1)}\}$. Now, the prior probabilities and remaining input state energy $(1-t_1)|\alpha|^2$ are by construction such that $P_I = 2(\epsilon' - (C')^2)/(1-2C')$ with $(C')^2 = \epsilon'(K')^2$, $(K')^2 = |\langle -\sqrt{1-t_1}\alpha|\sqrt{1-t_1}\alpha\rangle|^2$, and $\epsilon' = P_C^{(1)}\big(1 - P_C^{(1)}\big)$, corresponding to the so-called "single-state" domain [63, 111]. In this domain, the POVM element corresponding to the less likely input state is zero ($\hat{\Pi}_2 = 0$) and the other two POVM elements $\hat{\Pi}_1$ and $\hat{\Pi}_?$ are non-zero. Thus, the purpose of the second temporal mode is to implement a measurement which optimally distinguishes between $\hat{\Pi}_1$ and $\hat{\Pi}_?$ with minimal error. The outcome after the second temporal mode is then

either the inconclusive outcome with probability P_I or deciding the input state is $|\alpha\rangle$ with probability $P_C = P_C^{(s)}$. The quantity $P_C^{(s)}$ is the probability of correct detection in the single-state domain given the input states $\{|\pm\sqrt{1-t_1}\alpha\rangle\}$ with prior probabilities $\{P_C^{(1)}, 1-P_C^{(1)}\}$ [63]:

$$P_C^{(s)} = \frac{(C')^2\big(2\tilde{P}_C^{(1)}+1\big) + P_C^{(1)}\tilde{P}_C^{(1)}\big(2P_C^{(1)}-1\big) + 2\epsilon' K'\tilde{K}'\sqrt{(C')^2-(D')^2}}{\big(2P_C^{(1)}-1\big)^2 + 4(C')^2}, \quad (2.15)$$

with $(D')^2 = \big(P_C^{(1)} - P_I\big)\big(1-P_C^{(1)}-P_I\big)$, $\tilde{K}' = \sqrt{1-(K')^2}$, and $\tilde{P}_C^{(1)} = P_C^{(1)} - P_I$. The value of $P_C^{(s)}$ is such that $P_E = 1 - P_I - P_C^{(s)}$ is the smallest attainable error probability, thus achieving the optimally inconclusive measurement.

The measurement of the second temporal mode during the time interval $t_1 < t \leq 1$ is also implemented using displacement operations, photon counting, and feedback with a particular displacement waveform. A main result of Ref. [63] is that the Dolinar-*like* receiver which implements the required binary measurement in the second temporal mode has a displacement waveform of:

$$u_2(t) = \frac{(-1)^{N(t)+N_0}\alpha}{\sqrt{1-4v(1-v)K^{2t}}}. \quad (2.16)$$

This waveform implements a Dolinar receiver but with the prior probability p in Eq. (2.11) replaced by v. The quantity $N_2(t)$ ($N_2(0) = 0$) is the number of photons detected during $t_1 < t \leq 1$ and N_0 determines the phase of the LO at the beginning of the second temporal mode. If $v > 0.5$ then $N_0 = 0$ and $N_0 = 1$ otherwise where v is found by [63]:

$$v = \frac{1}{2}\left(1 + \frac{v_b}{v_a}\right), \quad (2.17)$$

$$v_a = \Big(P_C^{(1)}K'\Big)^2 - \Big(2P_C^{(s)} - P_C^{(1)}\Big)^2, \quad (2.18)$$

$$v_b = -\Big(P_C^{(1)}\tilde{K}'\Big)^2 + 2\Big(2P_C^{(s)} - P_C^{(1)}\Big)\frac{\tilde{K}'}{K'}\sqrt{P_C^{(s)}\Big(P_C^{(1)} - P_C^{(s)}\Big)}. \quad (2.19)$$

The measurement in the second temporal mode can be interpreted as a Dolinar receiver but the strategy is operating as if the two input states have prior probabilities $\{v, 1-v\}$ instead of $\{P_C^{(1)}, 1-P_C^{(1)}\}$ from Eq. (2.14). In a similar way to the standard

Dolinar receiver and Eq. (2.12), the temporal evolution for the probabilities of correct $P_C(t)$ and incorrect $P_E(t)$ discrimination during the second temporal mode where $t_1 < t \leq 1$ can be written as:

$$P_C(t+\delta t) = P_C(t)\Big(1 - |u_2(t) - \alpha|^2 \delta t\Big) + \big(P_C^{(1)} - P_C(t)\big)|u_2(t) + \alpha|^2 \delta t, \quad (2.20)$$

$$P_E(t+\delta t) = P_E(t)\Big(1 - |u_2(t) + \alpha|^2 \delta t\Big) + \big(1 - P_C^{(1)} - P_E(t)\big)|u_2(t) - \alpha|^2 \delta t, \quad (2.21)$$

where $P_I(t) = 1 - P_C(t) - P_E(t)$, $P_E(t_1) = 1 - P_C(t_1)$, and $P_C(t_1) = P_C^{(1)}$ if $v > 0.5$ or $P_C(t_1) = 0$ otherwise. The total measurement waveform for the displacement operation is constructed by combining the Dolinar waveform in Eq. (2.11) for the first temporal mode, and the Dolinar-*like* waveform in Eq. (2.16) for the second temporal mode:

$$u(t) = \begin{cases} u_1(t) = \frac{(-1)^{N_1(t)}\alpha}{\sqrt{1-4p(1-p)K^{2t}}} & 0 \leq t \leq t_1 \\ u_2(t-t_1) = \frac{(-1)^{N_2(t-t_1)+N_0}\alpha}{\sqrt{1-4v(1-v)K^{2(t-t_1)}}} & t_1 < t \leq 1. \end{cases} \quad (2.22)$$

Note that the waveform for the second mode is applied for the interval $t_1 < t \leq 1$ but the waveform is calculated using an interval $0 < t \leq 1 - t_1$. In this sense, the receiver implements a standard Dolinar receiver in the first temporal mode for the states $\{|-\sqrt{t_1}\alpha\rangle, \sqrt{t_1}\alpha\rangle\}$ up until $t = t_1$. Then, the receiver implements a Dolinar-*like* receiver in the second temporal mode for the states $\{|-\sqrt{1-t_1}\alpha\rangle, \sqrt{1-t_1}\alpha\rangle\}$ by operating as if input states have prior probabilities $\{v, 1-v\}$ from Eq. (2.17-2.19).

Figure 2.13(a) shows examples of the displacement magnitude $|u(t)|$ for an inconclusive probability $P_I = 0.25$ for $|\alpha|^2 = 0.2$, 0.4, and 0.6 in blue, orange, and yellow, respectively. The jumps in the waveform occur at the switching time $t_1 \approx$0.70, 0.57, and 0.42, for each investigated mean photon number $|\alpha|^2$, respectively. Figure 2.13(b) shows an example of the time evolution of the displacement waveform $u(t)$ for $|\alpha|^2 = 0.2$ and $P_I = 0.25$. Initially, the receiver implements a Dolinar receiver using the waveform $u_1(t)$ from Eq. (2.11) to discriminate between the two input states with minimum error. The sign of $u(t)$ (phase of the LO) is switched every time a photon is

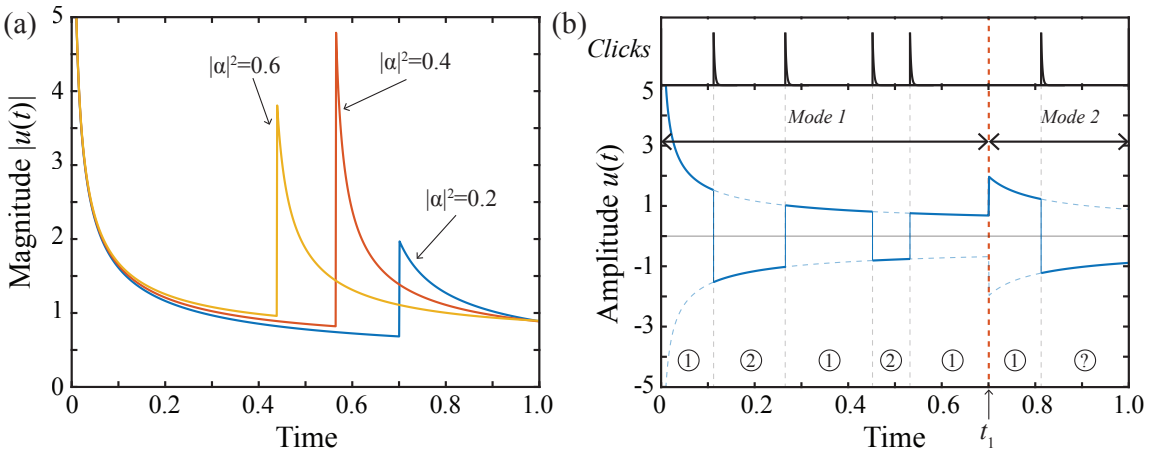

Figure 2.13: (a) Displacement magnitude $|u(t)|$ which implement an optimally inconclusive measurement for $|\alpha|^2$=0.2, 0.4, and 0.6 in blue, red, and yellow, respectively. The jumps in magnitude correspond to the switching time t_1 where the measurement switches from the first temporal mode to the second temporal mode. (b) Example of the time evolution of the displacement waveform $u(t)$ when the receiver detects photons (clicks) at different times. Every time a photon is detected, the LO phase (sign of $u(t)$) switches along with the current hypothesis (circles) for the input state. Prior to t_1, the hypothesis switches between the two input states each detection and after t_1, switches between the best guess after the first mode (state 1 as shown) and the inconclusive result.

detected (top panel). The current hypothesis (circled numbers) for the input state is given by the number of detected photons so far. After $t_1 \approx 0.7$ with $N_0 = 0$, the receiver implements the Dolinar-*like* receiver using the waveform $u_2(t)$ from Eq. (2.16) to discriminate between the most likely state after the first temporal mode and the inconclusive outcome. The overall outcome of the measurement is either correctly guessing the state with probability $P_C^{(s)}$ or an inconclusive result with probability P_I and the probability of error is: $P_E = 1 - P_I - P_C^{(s)}$.

2.3.2 Experimental Implementation

We use the same experimental setup shown in Fig. 2.3 to implement the optimally inconclusive measurement. We obtain the experimental displacement waveform $u(t)$ by replacing α in Eq. (2.11) and (2.16) with $\sqrt{\eta}\alpha$, where our system efficiency is $\eta = 0.72$. Although the measurement calls for extremely large values of the ratio

between the displacement energy and input energy $R = |u(t)|^2/|\alpha|^2$, experimentally there is a maximum ratio R which we can implement. Due to the extinction ratio ($\approx$ 20dB) of the amplitude modulator in the LO arm of the interferometer, we set the maximum of this ratio to $R = 50$ (See Appendix B.3 for details).

We discretize time into 1024 time bins of 160 ns each ($\approx$ 164 μs total) where a photon can potentially be detected to approximate a continuous measurement in our demonstration. This choice of time bin size allows our experiment to be implemented at a reasonable repetition rate of 4 kHz while also allowing the FPGA to easily and reliably transfer the entire measurement record for each time bin to the computer for processing. The magnitude of the displacement waveform $|u(t)|$ for a specific $|\alpha|^2$ and P_I is stored in a pre-calculated look-up table (LUT) and evolves according Eq. (2.22). The phase of the LO switches between 0 and π every time a photon is detected such that $u(t)$ switches between $|u(t)|$ and $-|u(t)|$. The phase of the LO at time $t = t_1$ is determined by N_0, which is given by the value of v in Eq. (2.17). The FPGA registers all photon detections and saves the result for all time bins such that the experimental measurement probabilities $\{P_E^{exp}(t), P_C^{exp}(t), P_I^{exp}(t)\}$ can be readily reconstructed as a function of time in post-processing with MATLAB.

Figure 2.14(a) shows an example of the time evolution of $P_E(t)$ (solid blue), $P_C(t)$ (solid orange), and $P_I(t)$ (solid yellow) as the measurement progresses for $|\alpha|^2 = 0.2$ and $P_I \approx 0.31$. The solid lines show the theoretical expectations (Eqs. (2.12), (2.20), and (2.21)) and the points show the experimental results for $P_E^{exp}(t)$, $P_C^{exp}(t)$, and $P_I^{exp}(t)$ every 50 time bins averaged across five different experimental runs. Figure 2.14(b) shows a zoom into the region corresponding to approximately $0.5 < t \leq 1$ and $0.05 < P < 0.35$ which contains the evolution of the inconclusive probability $P_I(t)$ (solid yellow). The measurement switches from the first mode to the second mode at $t_1 \approx 0.57$. The points show the average experimental probabilities with error bars corresponding to one standard deviation. The dashed gray line shows the ideal Helstrom bound ($\eta = 1.0$) for MESD with $|\alpha|^2 = 0.2$, which the receiver surpasses at $t \approx 0.7$. We note that this is possible at the expense of introducing inconclusive

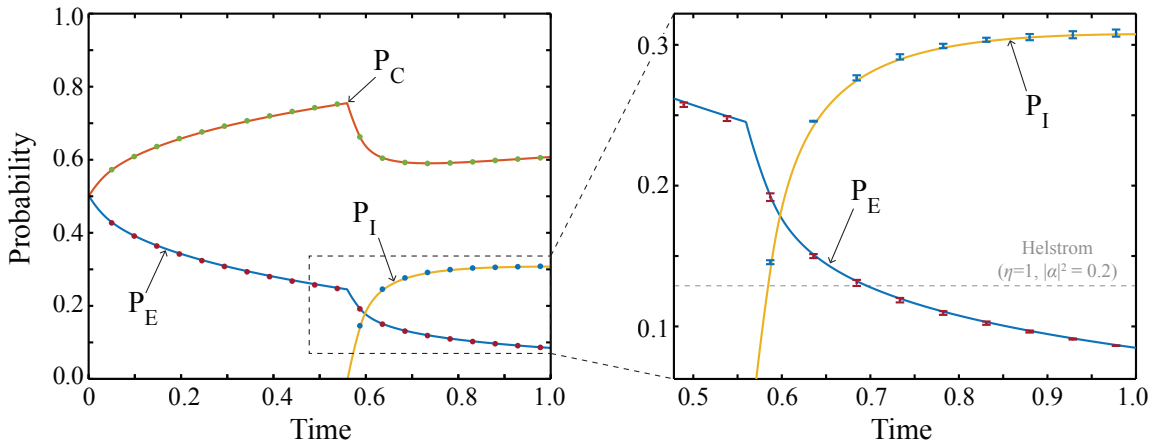

Figure 2.14: Time evolution of the correct probability (P_C, red), error probability (P_E, blue), and inconclusive probability (P_I, yellow) for a strategy where $|\alpha|^2$=0.2 and $P_I \approx$0.31. Solid lines show the theoretical expectations based on Monte Carlo simulations including experimental imperfections and the points with error bars represent the experimentally measured probabilities which show excellent agreement.

events which contribute to P_I in the optimally inconclusive measurement.

Taking the final measured inconclusive probability $P_I^{exp}(t=1)$ and error probability $P_E^{exp}(t=1)$, we obtain a single experimental point at the coordinates $(P_I^{exp}(1), P_E^{exp}(1))$ for a given input mean photon number $|\alpha|^2$. Figure 2.15(a) shows our total experimental results for the optimally inconclusive measurement for $|\alpha|^2 = 0.2$, 0.4, and 0.6 in blue, orange, and yellow, respectively. The data points show the experimental results $(P_I^{exp}(1), P_E^{exp}(1))$ and the error bars represent one standard deviation across five experimental runs each with 5×10^4 individual experiments. The black lines show the theoretical expectations from Monte Carlo simulations of the experiment incorporating experimental imperfections such as dark counts, visibility, and finite value of $R = 50$. We note that for the data points, the inconclusive probability is the experimentally obtained $P_I^{exp}(1)$, not the theoretically expected value P_I. The dashed gray lines show the corresponding Helstrom error probability for an ideal measurement ($\eta = 1$) for each input mean photon number. For these three values of $|\alpha|^2$, the uncorrected ($\eta = 0.72$) optimally inconclusive measurement can reach errors below the ideal Helstrom bound when $P_I > 0.15$, i.e. the point for $|\alpha|^2 = 0.6$ at (0.2, 0.02) represents the imperfect

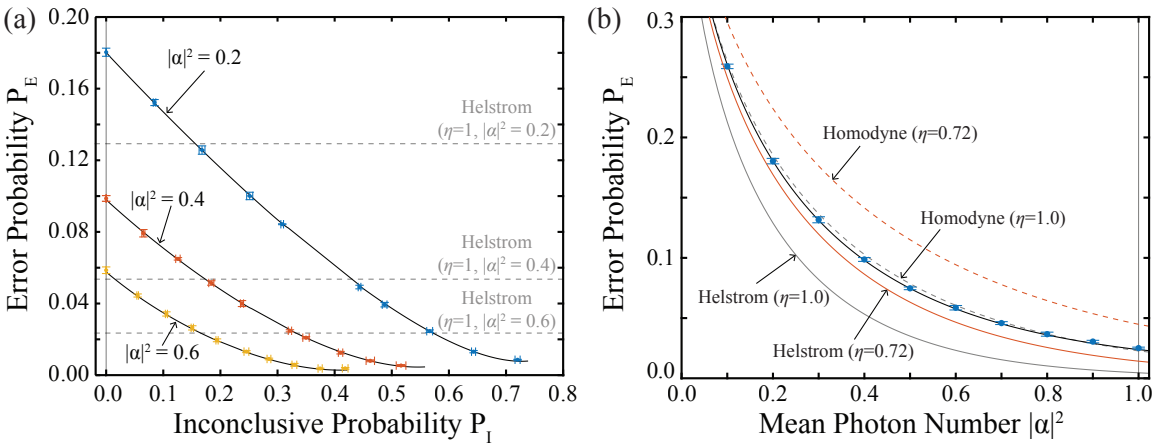

Figure 2.15: (a) Experimental results (points) for the optimally inconclusive measurement of equiprobable BPSK coherent states for $|\alpha|^2$=0.2, 0.4, and 0.6, in blue, orange, and yellow, respectively. Each point corresponds to a the experimentally measured value of $P_I^{exp}(1)$ and $P_E^{exp}(t)$ and the error bars represent the one standard deviation across five runs of 5×10^4 independent experiments each. The solid lines show the expected results including experimental imperfections and the dashed gray lines show the ideal ($\eta = 1.0$) Helstrom bound for each $|\alpha|^2$. The receiver achieves an error probability below the ideal Helstrom bound with an imperfect experimental implementation by allowing for inconclusive results. (b) Experimental results (blue points) for the Dolinar receiver for equiprobable BPSK coherent states. The gray and red solid lines show the Helstrom bound for $\eta = 1.0$ and $\eta = 0.72$ and the dashed lines show the error for a homodyne receiver.

experimental demonstration surpassing the ideal Helstrom bound.

The above framework from Ref. [63] for the optimally inconclusive receiver can also be used to implement a MESD measurement ($P_I = 0$) with a Dolinar receiver. In Ref. [53], a Dolinar receiver for OOK coherent states was implemented and achieved the equivalent Helstrom bound. However, the optimal encoding for binary coherent states for a fixed average energy $|\alpha|^2$ is a BPSK alphabet [69] due to having the smallest overlap ($|\langle -\alpha|\alpha\rangle|^2 < |\langle 0|\sqrt{2}\alpha\rangle|^2$), and thus the highest distinguishability. We demonstrate the optimal MESD receiver for the optimal binary encoding by implementing a Dolinar receiver for BPSK coherent states. Figure 2.15(b) shows the experimental results (blue points) and the expected error probability (solid black) as well as the corresponding Helstrom (solid) and homodyne limits (dashed) for an ideal experiment with $\eta = 1.0$ (gray) and for our overall efficiency of $\eta = 0.72$ (red). The

deviation of the data from the equivalent Helstrom bound is mainly due to detector dark counts and the finite value of $R = 50$. A MESD measurement in Fig. 2.15(b) corresponds to the y-axis in Fig. 2.15(a) and a USD measurement corresponds to the x-axis.

Similar to a Dolinar receiver, this approach can also implement an optimal unambiguous measurement for USD where $P_E = 0$ and $P_I = P_I^{USD}$. However, Fig 2.15(a) shows that experimental imperfections such as dark counts and non-ideal visibility prevent the receiver from achieving exactly $P_E = 0$. Nevertheless, the above framework allows for finding the measurement (i.e. displacement waveform) which in principle can implement an unambiguous measurement. While a MESD receiver can be implemented solely by $u(t) = u_1(t)$, a USD receiver is implemented by using both waveforms ($u_1(t)$ and $u_2(t)$) as opposed to just $u_2(t)$, as one may suspect. The receiver implements a Dolinar receiver in the first temporal mode, but attempts to displace the most likely state to vacuum ($|u(t)| = |\alpha|$) once in the single-state domain for the second temporal mode. Although the resulting P_I is the minimal achievable value and the same as the simple implementation in [96], it is interesting to note that this USD measurement can also be implemented with the above Dolinar-*like* strategy.

2.3.3 Discussion

Optimally inconclusive measurements represent a more general class of quantum measurements which include the MESD and USD paradigms. An interpretation of these optimally inconclusive measurements is that by definition they have an optimal post-selection threshold to achieve the minimal possible error for a given inconclusive probability. Thus, an optimally inconclusive measurement can potentially be used for quantum key distribution and compared to other BPSK-based protocols [31, 98, 115] with post-selection. These measurements for binary encoded coherent states can be implemented with linear optics, photon counting, and fast feedback. Although theoretically this is difficult to extend to higher dimensional spaces, understanding

the full capabilities of these three measurement tools is certainly a worthwhile endeavor which may require experimental and theoretical efforts to push our understanding of non-Gaussian measurements. Nevertheless, this result demonstrates that implementable Dolinar-like receivers can be used to perform a wide variety of projective and non-projective measurements within a two dimensional space.

2.4 Conclusion

In this chapter, we addressed three important problems in the area of binary state discrimination using non-Gaussian measurements: (1) robustness to measurement imperfections, (2) optimized communication strategies for noisy channels, and (3) generalized measurements beyond MESD and USD. Robustness to experimental imperfections can be acquired by allowing the receiver to have photon number resolving capabilities, which also factors into the optimization of non-Gaussian measurements. The photon number resolution mainly provides robustness to non-ideal visibility due to experimental imperfections as well as also detector dark counts and after-pulsing. In situations with communication over noisy communication channels, optimization of the coherent state alphabet jointly with the receiver provides the best combination of robustness and sensitivity. Such optimized communication strategies can improve information transfer over across a phase diffusion channel. Interestingly, adjusting the alphabet can also account for imperfections in the receiver in a similar way to PNR. Finally, more general measurements in two dimensional spaces can be implemented using displacements, photon counting, and feedback to interpolate between MESD and USD type strategies. This shows that adaptive photon counting has the capability to implement very general POVMs which can potentially be applied to higher dimensional discrimination problems.

Chapter 3

Discriminating Four Coherent States

Increasing the size of the coherent state alphabet used for communication will increase the achievable amount of information transfer. Using an alphabet consisting of four coherent states allows a sender to encode two bits of information instead of a single bit using two states. In terms of distinguishability, the most energy efficient alphabet for four coherent states is one where the states have the same mean photon number but with different phases, referred to as quaternary phase-shift-keyed (QPSK) coherent states [50, 58, 116]. The conventional measurement strategy for QPSK states is a heterodyne receiver, where both quadratures of the input field are measured, and the resulting error probability is bounded by QNL_{QPSK} [5]. However, non-Gaussian strategies using displacement operations and single photon counting can surpass the QNL_{QPSK} [51, 117, 118].

A benefit of non-Gaussian measurements with simple, non-adaptive implementations is that they can be operated at high bandwidth compared to adaptive strategies. In addition, photon-counting based receivers that are not adaptive require minimal resources in terms of space and complexity and can more readily provide near term advantages for realistic optical communication [118–123].

By utilizing the smallest number of individual hypothesis testing stages, on/off detection, and inherently stable experimental setups, the capabilities of such non-Gaussian receiver strategies can be explored [82, 118]. On the other hand, adaptive strategies based on signal nulling for hypothesis testing have been proposed [124–130] and demonstrated unconditionally surpassing the QNL_{QPSK} at relatively high input mean photon numbers when using PNR detection [51, 54, 117]. However, the performance of these receivers is worse than the QNL_{QPSK} in the single photon regime. Measurements which can reach below the QNL_{QPSK} in the low-power limit have potential applications in quantum communication [25–27, 30, 31, 131, 132], computation [34–37, 133], and networking [19, 134]. Similar to the receivers for BPSK coherent states studied in Ch. 2, a key component for achieving sub-QNL_{QPSK} performance in the single photon regime is optimization of the displacement operation [82, 83, 118, 119]. Although, now there are different classes of optimization for adaptive strategies, each with different levels of complexity given the nature of feedback-based receivers. Optimized strategies are critical to low power communications and a fundamental problem is finding receivers which can still be implemented in the presence of non-trivial or unknown channel noise [18, 71–73, 75, 95, 101, 102, 135]. Optimization of the receiver strategies fundamentally requires knowledge of the photon counting likelihood functions and therefore a mathematical description of the communication channel. Thus, alternative approaches which can be used to implement and optimize sub-QNL_{QPSK} measurements without a known channel description are of great interest.

In this chapter, we discuss experimental and theoretical investigations which address these problems of scalable, high-bandwidth receivers [82], optimized low-power measurements [83], and strategies which can be implemented in the presence of unknown channel noise. In Section 3.1, we explore the implementation of a non-adaptive, minimal resource strategy and identify the critical parameters which impact its performance. In section 3.2, we investigate the capabilities of different optimization strategies for an adaptive receiver in the single photon regime. In Section 3.3 we

utilize reinforcement learning concepts to develop alternative approaches to optimized measurements, which can potentially be implemented in the presence of unknown channel noise.

3.1 Optimized Single-Shot QPSK Receiver

Photon counting based measurements which can readily scale to current communication bandwidths (GHz) provide a path towards real implementations of sub-QNL receivers. In addition, non-adaptive strategies which do not require feedback and use the minimum amount of "resources" are more directly applicable to current communication systems. One such strategy proposed in Ref. [119] for QPSK states uses a small number of fixed, optimized displacement operations followed by single photon detection to implement multiple hypothesis testing simultaneously instead of sequentially. The energy distribution for each hypothesis test and displacement magnitudes are also optimized such that the error probability is minimized. We implement this measurement using an inherently stable polarization interferometer which does not require active phase stabilization [82]. We also extend the theoretical analysis from Ref. [119] to include experimental imperfections such as interference visibility and dark counts and identify the limitations of these strategies [82].

3.1.1 Measurement Strategy

Figure 3.1(a) shows the non-adaptive QPSK discrimination strategy proposed in Ref. [119]. The input state $|\alpha_k\rangle = |\alpha e^{ik\pi/2}\rangle$ with $k \in \{0, 1, 2, 3\}$ is first split into three detection arms. Each arm is set to test a different hypothesis for the input state using displacements $\hat{D}(\beta_i)$ and single photon detection. The phase of the displacement $\arg(\beta_i)$ in arm i is set equal to the phase of the input state being tested and the magnitude $|\beta_i|$ is optimized such that the error probability is minimized. The displaced states $\hat{D}(\beta_i)|\alpha_k\rangle = |\alpha_k - \beta_i\rangle$ are then each detected by a single photon

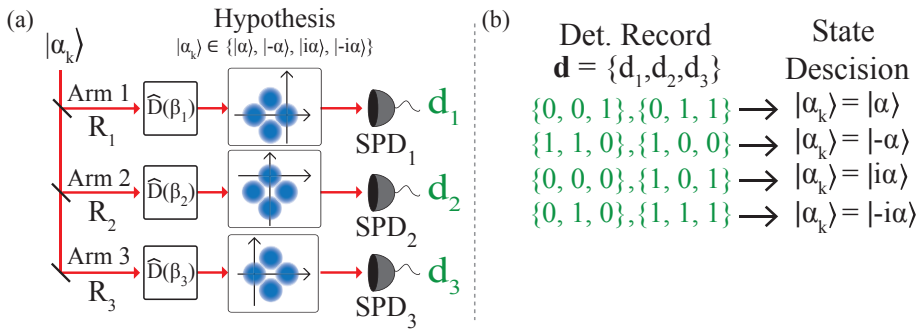

Figure 3.1: (a) Diagram of the optimized single-shot discrimination strategy. The input state $|\alpha_k\rangle$ is split into three detection arms with splitting ratios $\{R_1, R_2, R_3\}$. Each arm tests a different hypothesis for the input state by first displacing by an optimal value $\hat{D}(\beta_i)$ followed by single photon detection. (b) Based on the total detection record $\mathbf{d} = \{d_1, d_2, d_3\}$, the decision for the input is the state with maximum *a posteriori* probability. Figure adapted with permission from [82] © The Optical Society.

detector (SPD) with PNR(1), which yields a complete detection of $\mathbf{d} = \{d_1, d_2, d_3\}$ where d_i is the detection in the i^{th} arm. The posterior probability distribution $P(\alpha_k|\mathbf{d})$ given the measurement outcome $\mathbf{d}$ can be used to infer the input state and is given by: $P(\alpha_k|\mathbf{d}) = P(\alpha_k|d_1, \beta_1)P(\alpha_k|d_2, \beta_2)P(\alpha_k|d_3, \beta_3)$. The posterior probability depends on each displacement $\boldsymbol{\beta} = \{\beta_1, \beta_2, \beta_3\}$ as well as the fraction of energy $R_k|\alpha|^2$ ($\mathbf{R} = \{R_1, R_2, R_3\}$) sent to each arm with $R_1 + R_2 + R_3 = 1$. The overall error probability is then given by:

$$P_E = 1 - \frac{1}{4}\sum_{k=1}^{4} P(\alpha_k|\alpha_k, \boldsymbol{\beta}, \mathbf{R}), \tag{3.1}$$

where $P(\alpha_k|\alpha_k, \boldsymbol{\beta}, \mathbf{R})$ is the probability of correct discrimination. This error probability can then be minimized with respect to $\boldsymbol{\beta}$ and $\mathbf{R}$ for a fixed $|\alpha|^2$. Figure 3.1(b) shows the possible detection outcomes $\mathbf{d} = \{d_1, d_2, d_3\}$ and the corresponding decision for the input state. There are in total $2^3 = 8$ possible detection outcomes $\mathbf{d}$ since there are three detection arms each with a PNR(1) detector. The decision for the input is the state with maximum posterior probability given the detection $\mathbf{d}$), resulting in two different detection outcomes corresponding to each input state.

Figure 3.2(a) shows the expected error probability for different combinations of experimental parameters. The colored lines show the error probability for different

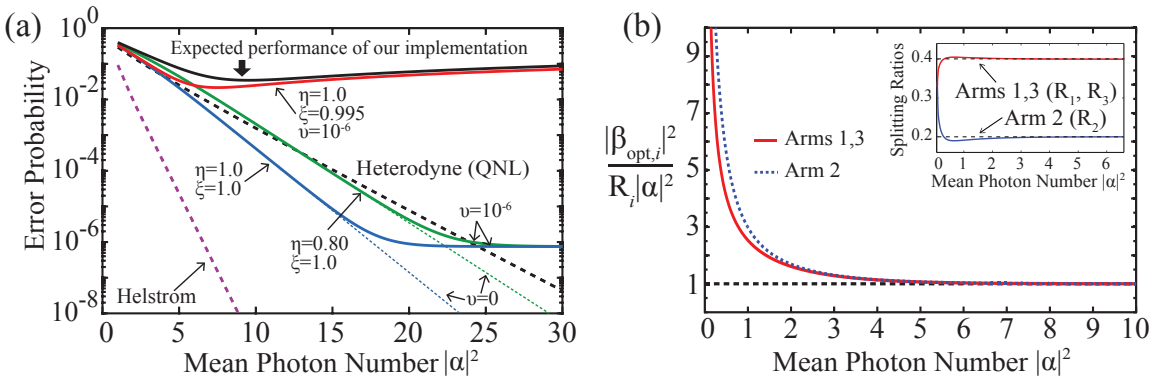

Figure 3.2: (a) Expected results for the strategy for different combinations of experimental parameters. The red, greed, and blue lines show the error probability for $(\eta, \xi) = (1.0, 0.995)$, (0.80, 1.0), and (1.0, 1.0), respectively, with a dark count rate of $\nu = 0$ (dashed) and $\nu = 10^{-6}$ (solid). The purple and black dashed lines show the Helstrom bound and QNL, while the black solid line shows the expected performance for our experimental implementation. (b) Optimal displacement ratios for each detection arm as a function of input mean photon number $|\alpha|^2$. The inset shows the optimal splitting ratios $\mathbf{R}$=$\{R_1, R_2, R_3\}$. Figure adapted with permission from [82] © The Optical Society.

pairs of (η, ξ) with (1.0, 0.995) in red, (0.80, 1.0) in green, and (1.0, 1.0) in blue, respectively, both without (dashed) and with (solid) dark counts of $\nu = 10^{-6}$. These cases are calculated using fixed splitting ratios of $\mathbf{R} = \{0.40, 0.20, 0.40\}$. The black solid line shows the expected error for our implementation given experimental imperfections, and the black dashed line shows the QNL_{QPSK} which is given by an ideal heterodyne receiver. Figure 3.2(b) shows the optimal displacement ratios $|\beta_{opt,i}|^2/R_i|\alpha|^2$ corresponding to the red line in (a) where $\eta = 1.0$ and $\xi = 0.995$ for each detection arm. The optimal displacement ratio for each arm asymptotes to a ratio of $|\beta_{opt,i}|^2/R_i|\alpha|^2 = 1$ as the input power increases and is slightly higher for Arm 2, which tests the hypothesis $|\alpha_k\rangle = |i\alpha\rangle$. The inset shows the optimal splitting ratios for each arm, which is initially $\mathbf{R} = \{0.33, 0.33, 0.33\}$ but asymptotes to $\mathbf{R} = \{0.40, 0.20, 0.40\}$. We think that the optimal splitting ratio for Arm 2 converges to 20% of the total power due to the fact that the state which Arm 2 is testing for ($|i\alpha\rangle$) has the most overlap between the other two states, i.e. $|\langle i\alpha| \pm \alpha\rangle|^2 > |\langle \pm\alpha| \mp \alpha\rangle|^2$.

3.1.2 Experimental Setup

We implement the measurement strategy described above using an inherently stable polarization interferometer with three separate detection arms [82]. Figure 3.3 shows the experimental setup used to perform QPSK state discrimination. Optical pulses are generated using a continuous Helium-Neon (HeNe) laser and a pulsed acousto-optic modulator (AOM) followed by an attenuator. The input state is encoded into the vertical polarization component of the field and the LO is encoded into the horizontal component. The input state and LO co-propagate through the interferometer and are split into the three detection arms with optimized splitting by a 60/40 BS and a tilted 50/50 BS to achieve the desired splitting ratios **R**. We implement the displacement operation using a half-wave plate (HWP) in each arm to slightly rotate the polarization

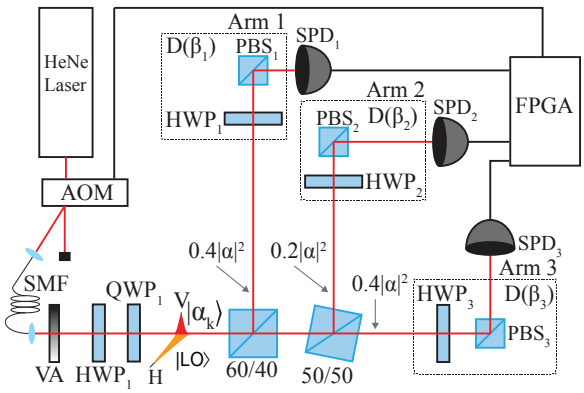

	Arm 1	Arm 2	Arm 3
DE	0.869(2)	0.731(3)	0.832(2)
Visibility	0.991(1)	0.990(1)	0.993(1)
Dark Counts	$5.42(2)\times10^{-5}$	$1.22(4)\times10^{-3}$	$1.65(3)\times10^{-3}$
Hypothesis	$\|\alpha_k\rangle = \|\alpha\rangle$	$\|\alpha_k\rangle = \|i\alpha\rangle$	$\|\alpha_k\rangle = \|-\alpha\rangle$
PBS ER	$2.03(2)\times10^{-4}$	$3.13(2)\times10^{-4}$	$1.45(1)\times10^{-4}$

Figure 3.3: Polarization based interferometer setup. The input state and LO are prepared in the vertical and horizontal polarization, respectively, with a variable attenuator (VA), HWP, and QWP and are then split into the three detection arms. In each arm, the optimal displacement is implemented by rotating the polarization with a HWP followed interference at a polarizing beam splitter (PBS). The displaced state is then measured by a single photon detector and the experimental parameters for each detection arm are shown below the diagram. Figure adapted with permission from [82] © The Optical Society.

such that when the input and LO fields interfere on a polarizing beam splitter (PBS), the input is displaced in phase space by an optimal amount $\hat{D}(\beta_i)$. The state is then detected using a single photon detector with PNR(1) and the data is collected by a field programmable gate array (FPGA). The table below the diagram shows the experimental parameters of detection efficiency, visibility, dark counts, and extinction ratio for each arm.

3.1.3 Results

Figure 3.4 shows the experimental results of the optimized single-shot measurement for QPSK state discrimination. The red solid line shows the expected performance from Monte Carlo simulations and the blue points show the measured data with error bars corresponding to one standard deviation across five runs of the experiment. The green, black, and gray lines show the Helstrom bound, ideal heterodyne limit, and the error of a heterodyne receiver adjusted by our average detection efficiency of $\eta = 0.778$. We find that the experimental results lie slightly above the adjusted heterodyne limit

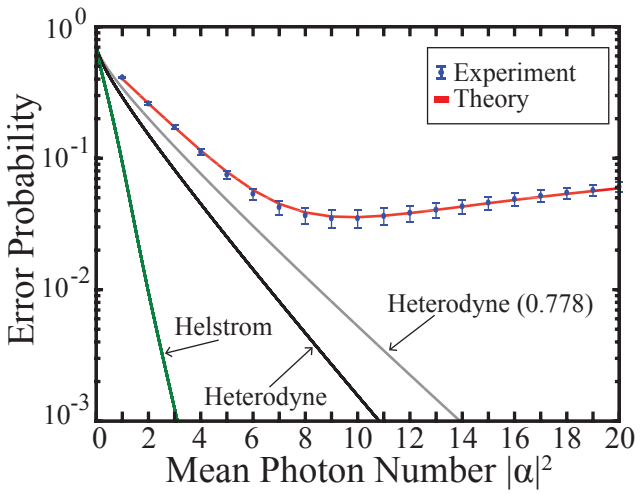

Figure 3.4: Measurement results showing the theoretical prediction (red solid) and experimental data (blue points). Error bars show the average across five runs of the experiment. The black (gray) lines show the ideal (adjusted) limit of a heterodyne receiver. Figure adapted with permission from [82] © The Optical Society.

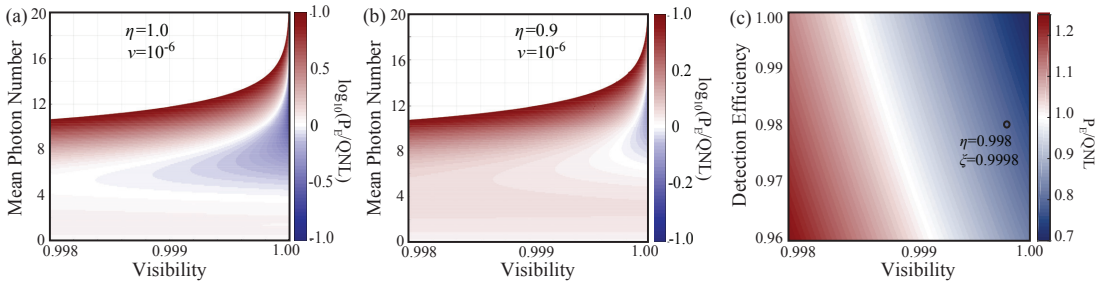

Figure 3.5: Error probability P_E of the simple non-Gaussian strategy compared to QNL_{QPSK} on a log scale as a function of the visibility and mean photon number for (a) $\eta = 1$ and (b) $\eta = 0.9$ with dark counts $\nu = 10^{-6}$. (c) Error probability ratio P_E/QNL_{QPSK} as a function of visibility ξ and detection efficiency η for $|\alpha|^2 = 6.0$. The black circle represents the capabilities of the strategy with a state-of-the-art detector ($\eta = 0.998$, $\nu \approx 0$) and beam splitters which achieve $\xi = 0.9998$. Figure adapted with permission from [82] © The Optical Society.

and closely match the theoretical predictions, which validates the theoretical proposal in Ref. [119]. This performance is mainly due to the non-ideal detection efficiency and interference visibility in each detection arm, but state-of-the-art components can be used to improve the performance of the receiver.

Figures 3.5(a,b) show the expected ratio of the error probability for the non-Gaussian receiver to a heterodyne receiver $\log_{10}(P_E/\mathrm{QNL}_{QPSK})$, as a function of the visibility ξ and input mean photon number $|\alpha|^2$ for (a) $\eta = 1.0$ and (b) $\eta = 0.9$ both with $\nu = 10^{-6}$. The blue regions correspond to when the non-Gaussian receiver shows an improvement over a heterodyne measurement ($P_E < \mathrm{QNL}_{QPSK}$). Figure 3.5(c) shows P_E/QNL_{QPSK} as a function of visibility ξ and detection efficiency η for a fixed mean photon number $|\alpha|^2 = 6.0$ with $\nu = 10^{-6}$. The blue area corresponds to the region in the experimental parameter space where implementations of these "minimum resource" measurements that don't require feedback operations provide an advantage over the QNL. As a specific example, the black point shows the benefit provided by a state-of-the-art superconducting detector with $\eta \geq 0.98$ and $\nu \approx 0$ [136] with high extinction ratio beam splitters which potentially allow for $\xi \geq 0.9998$ [137]. This point in the parameter space achieves an expected advantage of $P_E/\mathrm{QNL}_{QPSK} \approx 0.80$ at $|\alpha|^2 = 6.0$ using current technologies.

3.1.4 Discussion

This experimental demonstration validates the proposal from Ref. [119] and furthers our understanding of the limitations and capabilities in measurements which use "minimum resources". A minimum resource measurement is one which is as simple as possible in terms of complexity, i.e. no feedback, PNR(1), and a small number of simultaneous measurements. These types of receivers can potentially provide a near-term benefit over the QNL_{QPSK} due to their simplicity and ability to operate at high bandwidths. Furthermore, advances in high efficiency, low dark-count detectors and high quality optical components potentially allow for surpassing the ideal QNL_{QPSK} with a realistic implementation. We also note that recent theoretical and experimental work in Ref. [118] shows that further benefits can be obtained by also optimizing the phase of the displacements in addition to the magnitude.

3.2 Global Optimizations in the Low-Power Limit

Measurement receivers which achieve the best possible error probability through optimized and adaptive methods are critical to understanding the fundamental limits of non-Gaussian measurements. The limits of non-Gaussian receivers are especially important in the single-photon regime, where measurements need to be made on very faint pulses of light. Due to the adaptive nature of previously demonstrated strategies [51, 117], optimization of these receivers results in different levels of complexity depending on exactly what parameters are optimized. On one hand, the displacement magnitude for the entire measurement (a single parameter) can be optimized to minimize the error probability. On the other hand, the displacement for every possible measurement record can be optimized in a global way to obtain minimal error probability, but at the cost of exponentially many parameters. These optimizations allow adaptive non-Gaussian measurements to surpass the QNL in the single-photon regime, where communications are extremely power-limited.

3.2.1 Optimization Strategy

We study the performance of optimized versions of the adaptive non-Gaussian receiver demonstrated in Ref. [51] for QPSK coherent states. The adaptive receiver in Ref. [51] follows the Bayesian hypothesis testing strategy described in App. A with $L = 10$ adaptive measurement steps without photon number resolution, i.e. PNR(1). The optimization of this photon counting measurement can be split into three different strategies of increasing complexity for given values of L and PNR. The simplest method is one where the displacement magnitude $|\beta_j|$ for all adaptive steps $j \in \{1...L\}$ is optimized such that $|\beta_j| = |\beta_{opt}|$, corresponding to a single-parameter optimization and we refer to this approach as a "flat" optimization. A more complex strategy which improves the performance is obtained by optimizing the displacement magnitude for each adaptive step j. This will yield L optimal magnitudes such that $|\beta_j| = |\beta_{j,opt}|$ and is denoted as a "sequential" optimization. The best performance is achieved by a global or "historical" optimization, where the displacement magnitude is optimized for every possible measurement record. In this "historical" optimization strategy the displacement magnitude for all 2^L ($(m+1)^L$ for PNR(m)) detection histories needs to be optimized, which yields different LO magnitudes at each adaptive step conditioned on the particular measurement outcome up to that step. For example, the magnitude of the third adaptive step ($j = 3$) is dependent on the detections in the first two steps $\mathbf{d} = \{d_1, d_2\}$ such that for PNR(1), the optimal magnitude for the third step is equal to $|\beta_{opt}(\{d_1, d_2\})|$ and $|\beta_{opt}(\{0,0\})| \neq |\beta_{opt}(\{0,1\})| \neq |\beta_{opt}(\{1,0\})| \neq |\beta_{opt}(\{1,1\})|$. In this way, the LO magnitude for every possible detection history $\mathbf{d}$ at every adaptive step is optimized such that the receiver aims to achieve the smallest possible error probability.

Figure 3.6(a) shows the expected error probability for an ideal experiment ($\eta = 1.0, \xi = 1.0, \nu = 0$) with $L = 10$ adaptive steps. The brown, green, and blue solid lines show the error for the "flat", "sequential", and "historical" optimization strategies, respectively. The pink line shows the performance without optimization, which cannot surpass the QNL_{QPSK} (red) when $|\alpha|^2 \lessapprox 0.65$. On the other hand, all three

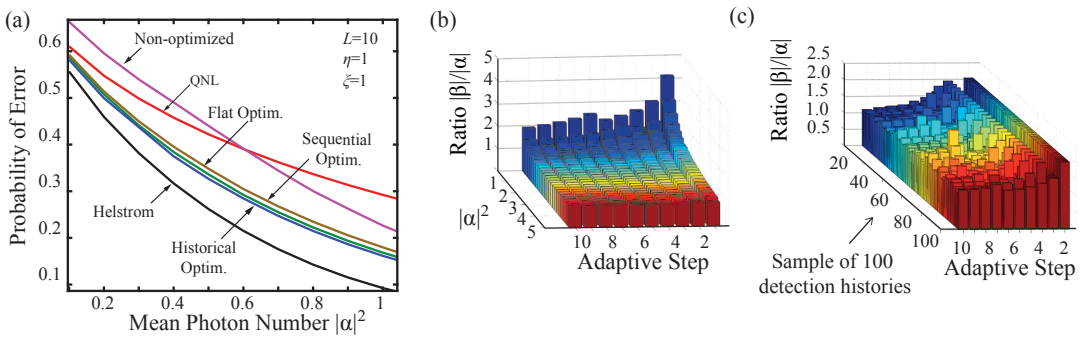

Figure 3.6: (a) Theoretical performance of different optimization strategies compared to the QNL (red) and Helstrom bound (black). The brown, green, and blue solid lines show the error for the "flat", "sequential", and "historical" optimization strategies, respectively, for ideal receiver ($\eta = 1.0, \xi = 1.0, \nu = 0$) with $L = 10$ adaptive steps. The purple line shows the result for a non-optimized measurement. (b) Optimal displacement ratios for each adaptive step for the "sequential" optimization strategy for different $|\alpha|^2$. (c) Sample of optimal displacement ratios for each adaptive step for the "historical" optimization strategy. Figure adapted from: *npj Quantum Information* 3, 43 (2017).

optimization strategies outperform the QNL_{QPSK} in this single photon regime. Figure 3.6(b) shows the optimal displacement ratios $|\beta_{j,opt}|/|\alpha|$ for each adaptive step in the "sequential" optimization at different values of $|\alpha|^2$. Figure 3.6(c) shows a sample of 100 possible detection histories and the corresponding optimal ratio $|\beta_{opt}(\mathbf{d})|/|\alpha|$ each adaptive step for the "historical" optimization strategy. The detection histories are enumerated by turning the binary detection record $\mathbf{d} = \{d_1, d_2...\}$ with $d_j \in \{0, 1\}$ into a decimal number.

3.2.2 Experimental Results & Discussion

We implement the optimized strategies using the experimental setup described in App. A.2 where the displacement is controlled by phase and amplitude modulators in the LO arm of the interferometer. A look-up-table (LUT) outputs the optimal displacement magnitude $|\beta_{opt}(\mathbf{d})|$ each adaptive step based on the current detection record $\mathbf{d}$. The optimization and LUT are pre-calculated in MATLAB and loaded into the FPGA, which implements the strategy. Figure 3.7 shows the experimental results for the optimized state discrimination strategies. The brown, green, and blue points

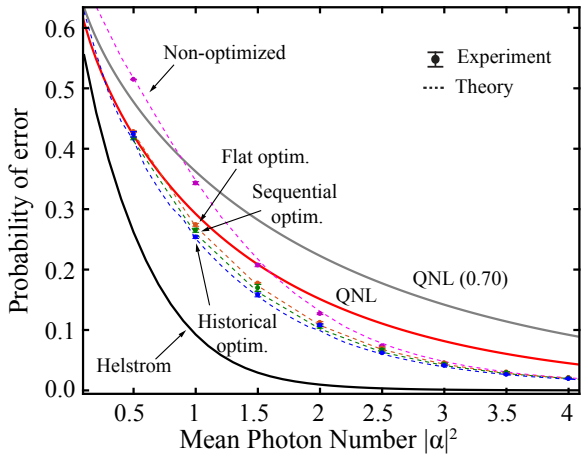

Figure 3.7: Experimental error probability (points) for the "flat" (orange), "sequential" (green), and "historical" (blue) optimization strategies and a non-optimized (pink) strategy. The black, red, and gray solid lines show the Helstrom bound, QNL_{QPSK}, and effective QNL_{QPSK} for our efficiency ($\eta = 0.70$), respectively. The dashed colored lines show the theoretically expected error from Monte Carlo simulations of the experiment. Figure adapted from: *npj Quantum Information* 3, 43 (2017).

show the results for the "flat", "sequential", and "historical" optimizations where the error bars correspond to the standard deviation across five experimental runs. While the non-optimized strategy (pink points) cannot surpass the ideal QNL_{QPSK} for $|\alpha|^2 \leq 1.5$, all of the optimization strategies do except for very small input powers. We also find that as the input energy increases beyond $|\alpha|^2 \approx 2.5$, the non-optimized error probability is almost equivalent to that of the optimized strategies.

This experimental demonstration shows that global optimizations can allow measurements based on adaptive photon counting to surpass the QNL_{QPSK} at low input powers. However, there are multiple different levels of optimizations a receiver can implement due to the adaptive nature of the measurement. Here, we find that more complex optimizations of the displacement magnitude can lead to more improvements over the QNL_{QPSK}. Measurements which surpass the QNL_{QPSK} in the single-photon regime are of particular importance to quantum communication protocols [25–27, 30, 31, 131, 132], which often rely on discrimination of coherent states with small powers. One interesting extension of this work will be to incorporate other experimental parameters into the optimization such as the duration of each adaptive

step (i.e. fraction of the total energy), the displacement phase, or the input alphabet, and to explore optimized measurements for higher order modulation formats.

3.3 Optimized Strategies for Unknown Channels with Reinforcement Learning

A basic assumption in constructing measurement strategies is that the receiver knows the properties of the communication channel, and therefore has a specific description of the states to be discriminated. For example, a pure loss channel simply re-scales of the magnitude of the states, and a phase noise channel rotates the states in a unitary way. Furthermore, the output states of a channel may not necessarily be pure states such as for phase diffusion channel, which can be described as a convex combination of states weighted by a Gaussian distribution in phase. The physical states that the receiver is attempting to measure are known for all of these channels in the sense that the likelihood function is given by: $\mathcal{L}(n|\mathcal{E}(\hat{\rho}_k),\beta) = \mathrm{Tr}\left[\mathcal{E}(\hat{\rho}_k)\hat{\Pi}_n(\beta)\right] = \langle n|D(\beta)\mathcal{E}(\hat{\rho}_k)D^\dagger(\beta)|n\rangle$, where $\mathcal{E}(\hat{\rho}_k)$ represents the output of the channel $\mathcal{E}(\cdot)$ when acting on the input state $\hat{\rho}_k$. Knowing the channel $\mathcal{E}(\cdot)$, and by extension the likelihood functions $\mathcal{L}(\cdot)$, allows for constructing measurement receivers based on Bayesian updating which can surpass the limits of conventional detection in the single photon regime using complex optimizations [83]. However, unknown channels pose a fundamental question: if one does not know or cannot calculate $\mathcal{E}(\hat{\rho}_k)$, then how can an optimized receiver be constructed?

Recently, studies have shown that certain measurement strategies can be found using reinforcement learning, where an agent (i.e. measurement receiver) attempts a particular strategy, obtains a reward based on performance, and then updates the strategy based on the reward [138, 139]. However, standard methods for finding solutions in a reinforcement learning framework either still rely on being able to calculate likelihood functions, or require learning an exponentially large table of values

for the displacement operations [138]. Thus, new approaches to finding optimized adaptive photon counting strategies are needed which don't depend on knowledge of the communication channel or an exponentially large table of values.

3.3.1 Alternative Representation of a Measurement

Adaptive photon counting measurements can be decomposed into two fundamental components. The first component implements the Bayesian updating using the likelihood functions to obtain the posterior distribution as described in Appendix A. The Bayesian aspect of the measurement is the receivers way of maintaining a belief (as a probability distribution) for the input states and updating that belief by incorporating new information gained from photon detections [51, 117]. If the receiver doesn't know the noise properties of the channel, and thus doesn't know the likelihood functions, then the standard Bayesian approach breaks down. Moreover, ideally the receiver implements an *optimized* measurement like the strategies discussed in Section 3.2, which also require knowledge of the likelihood functions. As an alternative to recursive Bayesian updating, the receiver can use an object which is capable of maintaining a belief for the input states (not necessarily as probabilities) and recursively updating that belief when new information is obtained through photon detections. Such an object or architecture ideally is inherently recursive and has a simple representation that is not dependent on the particular parameters of the problem at hand such as the number of adaptive steps L, or the PNR. A Long-Short Term Memory (LSTM) is a specific type of recurrent neural network from the field of machine learning which is well suited for this task [140–143]. An LSTM has the ability to maintain an internal memory over long time scales and update that memory every time step due to it's "gate" structure. In this work, we explore the use of an LSTM in optimized receivers based on adaptive photon counting as a method to replace recursive Bayesian updating used in previous strategies [83].

Figure 3.8(a) shows the architecture of the LSTM used to replace the Bayesian

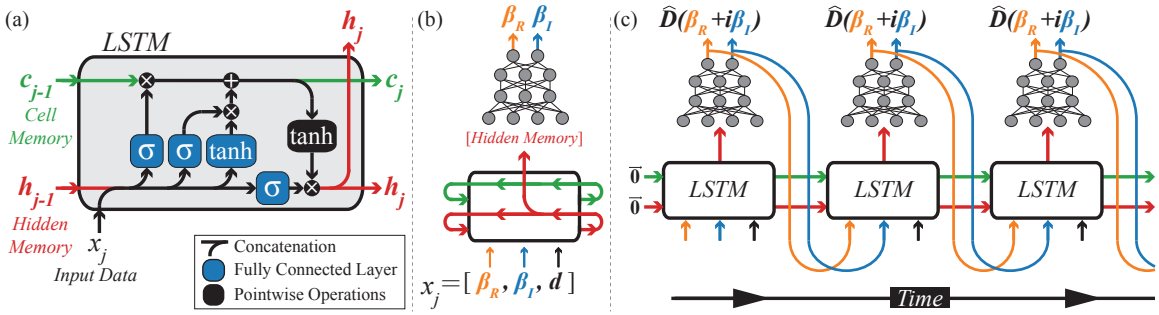

Figure 3.8: (a) Long-short term memory (LSTM) architecture which consists of three gates (input i_j, forget f_j, and output o_j) and two memories (cell C_j and hidden h_j). The memories maintain information from one adaptive step to another and the gates control how new information from photon detections are incorporated into the memories. (b) The hidden memory is passed to a deep neural network which outputs values for the real and imaginary parts of the displacement operation each adaptive step. (c) The LSTM can be temporally unrolled to show the flow of information between adaptive steps.

updating component of adaptive photon counting measurements for coherent states. The internal memory of the LSTM at adaptive step j is given by the cell memory C_j and a hidden memory h_j, which are both vectors. The three gates: forget f_j, input i_j, and output o_j, control how information flows from one adaptive step to the next. The gates also control how new information x_j is incorporated into the system and influences h_j, which is also the output of the LSTM. Merging arrows represent concatenation of data and the "σ" and "tanh" labels represent a single neural network layer followed by the application of the "sigmoid$(\cdot)$" or "tanh$(\cdot)$" activation functions. Thus, the LSTM is represented as a set of weights and biases which define each gate and memory in terms of the equations [140, 141, 144]:

$$i_j = \sigma(\mathbf{U}^{(i)}x_j + \mathbf{W}^{(i)}h_{j-1} + b^{(i)}), \tag{3.2}$$

$$f_j = \sigma(\mathbf{U}^{(f)}x_j + \mathbf{W}^{(f)}h_{j-1} + b^{(f)}), \tag{3.3}$$

$$o_j = \sigma(\mathbf{U}^{(o)}x_j + \mathbf{W}^{(o)}h_{j-1} + b^{(o)}), \tag{3.4}$$

$$\tilde{C}_j = \tanh(\mathbf{U}^{(c)}x_j + \mathbf{W}^{(c)}h_{j-1} + b^{(c)}), \tag{3.5}$$

$$C_j = \sigma(f_j \odot C_{j-1} + i_j \odot \tilde{C}_j), \tag{3.6}$$

$$h_j = \tanh(C_j) \odot o_j, \tag{3.7}$$

where $\mathbf{U}^{(l)}, \mathbf{W}^{(l)}$ and $b^{(l)}$ with $l \in \{i, f, o, c\}$ are trainable weight matrices and bias

vectors, respectively, and $\odot$ denotes element wise multiplication.

The second component of an adaptive measurement is a function which takes the current value of the memory h_j and outputs a magnitude $|\beta_j|$ and phase $\arg(\beta_j)$ for the displacement operation. Usually in adaptive measurement strategies, the receivers belief about the input states at adaptive step j is given by the prior probability distribution, which is simply the posterior distribution from adaptive step $j-1$ [51,87]. The displacement β_j is then obtained as: $|\beta_{j,opt}|e^{\hat{\phi}}$, where $\hat{\phi}$ is the phase of the state with the largest prior probability and $|\beta_{j,opt}|$ is the pre-calculated optimal magnitude as in Section 3.2 [83]. However, this approach still relies on being able to calculate the likelihood functions to obtain the posterior distribution each adaptive step and pre-calculate $|\beta_{j,opt}|$. A properly trained neural network (NN) can serve as a function *approximator* which in principle can map the current memory h_j on to an optimal value for the displacement without the need for likelihood functions. To this end, we construct a NN that is designed to find a new displacement value β_j for each adaptive step based on the current value of the LSTM memory h_j. A NN also has the benefit of being represented by a set of weight matrices and bias vectors which can be trained to achieve the best overall performance. Figure 3.8(b) shows the combined LSTM/NN architecture where the hidden h_j and cell C_j memories are fed-forward to maintain the memory of the network. The hidden memory h_j is passed to the NN which outputs two values corresponding to the real $\mathrm{Re}[\beta_j]$ and imaginary $\mathrm{Im}[\beta_j]$ parts of the displacement $\hat{D}(\beta_j)$ for the adaptive step. To obtain β_j for the current adaptive step j, the LSTM is evaluated with the input vector $x_j = \{\mathrm{Re}[\beta_{j-1}], \mathrm{Im}[\beta_{j-1}], 2d_{j-1}/m\}$ for a strategy with PNR(m). Figure 3.8(c) shows the temporally expanded LSTM/NN showing the flow of the information contained in the memories from one time-step to another.

This LSTM/NN architecture is represented by a set of numbers corresponding to every single individual weight and bias. In this work, the LSTM memories (h_j and C_j) are each 16 elements, corresponding to 1280 total parameters for the LSTM. We use a four layer network for the NN where the first three layers have 32 neurons and a "sigmoid" activation function, and the last layer has two outputs and a custom

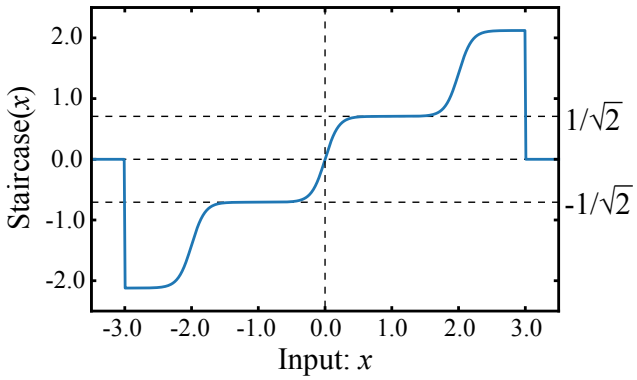

Figure 3.9: Custom staircase activation function which is tailored to a QPSK alphabet: $\{|\frac{|\alpha|}{\sqrt{2}}+i\frac{|\alpha|}{\sqrt{2}}\rangle, |-\frac{|\alpha|}{\sqrt{2}}+i\frac{|\alpha|}{\sqrt{2}}\rangle, |-\frac{|\alpha|}{\sqrt{2}}-i\frac{|\alpha|}{\sqrt{2}}\rangle, |\frac{|\alpha|}{\sqrt{2}}-i\frac{|\alpha|}{\sqrt{2}}\rangle\}$.

"staircase" activation function given by:

$$f(x)=\begin{cases}0 & x<-\frac{3w}{2}\\ 2h\Big(\frac{1}{2}\coth\big(\frac{a}{2}\big)\tanh\big(a(\frac{x}{w}-\lfloor\frac{x}{w}-\frac{1}{2}\rfloor-1)\big)+1+\lfloor\frac{x}{w}-\frac{1}{2}\rfloor\Big) & -\frac{3w}{2}\geq x\leq\frac{3w}{2}\\ 0 & x>\frac{3w}{2}\end{cases} \tag{3.8}$$

where $h=1/\sqrt{2}$ is the height of the "stairs", $w=2$ is the width, and $a=10$ controls how steep the transition is between steps. Figure 3.9 shows the "staircase" activation function for the output of the neural network. A single adaptive step j therefore comprised of three steps: (1) updating the LSTM memory by evaluating the LSTM with x_j, (2) evaluating the NN with the updated hidden memory h_j to obtain β_j, and (3) applying $\hat{D}(\beta_j)$ and detecting a particular number of photons d_j.

3.3.2 Replacing a Non-Optimized Strategy

We perform an initial study of receiver strategies based on the LSTM/NN architecture described above to replace non-optimized discrimination measurements. We train the LSTM/NN to implement a non-optimized, adaptive state discrimination strategy of the QPSK states $|\alpha_k\rangle=|\alpha e^{i\phi_k}\rangle$ with $\phi_k=\big(k+\frac{1}{2}\big)\frac{\pi}{2}$ and $k\in\{0,1,2,3\}$. This first study

allows for investigating and benchmarking how well the LSTM/NN architecture can implement a non-optimized measurement assuming a known channel. The typical adaptive Bayesian measurement performs hypothesis testing by adjusting the displacement β_j each adaptive step conditioned on photon detections [51,83], resulting in L detection results $\{d_j\}_L$ and corresponding displacement values $\{\beta_j\}_L$. Ideally, the LSTM/NN outputs displacement values close to $\{\beta_j\}_L$ when given the input sequence $\{d_j\}_L$, and we train the LSTM/NN to accomplish this. For example, if in the first few adaptive steps ($j = 1, 2, 3$) the displacement values and corresponding detections $\{\beta_j, d_j\}$ obtained from a Bayesian strategy are $\{\{\alpha e^{i\pi/4}, 2\}, \{\alpha e^{i3\pi/4}, 1\}, \{\alpha e^{i2\pi/4}, 0\}\}$, then given the fixed initial displacement $\beta_1 = \alpha e^{i\pi/4}$ and a detection $d_1 = 2$ the LSTM/NN should output the displacement $\hat{\beta}_2 \approx \alpha e^{i3\pi/4}$. Then, given a new detection of $d_2 = 1$ the LSTM/NN should output $\hat{\beta}_3 \approx \alpha e^{i2\pi/4}$, and so on. The displacement values $\beta_2 = \alpha e^{i3\pi/4}$ and $\beta_3 = \alpha e^{i2\pi/4}$ obtained from the Bayesian strategy are referred to as the target values because if the LSTM/NN output the target values exactly, then the LSTM/NN would exactly implement the Bayesian strategy, i.e. the target values are what we train the LSTM/NN to output.

We train the LSTM/NN using the gradient descent based Adam optimizer [145] with a learning rate of 10^{-4} for 200 epochs using the Keras framework [146] in Python. We use a mean-squared error cost function:

$$C = \sum_{i,j} \left(y_{j,R}^{(i)} - \hat{y}_{j,R}^{(i)}\right)^2 + \left(y_{j,I}^{(i)} - \hat{y}_{j,I}^{(i)}\right)^2, \tag{3.9}$$

where the target values $y_{j,R}^{(i)}$ and $y_{j,I}^{(i)}$ are the real and imaginary parts of the displacement $\beta_j^{(i)}$ in the j^{th} adaptive step of the recursive Bayesian strategy for the i^{th} sample, and $\hat{y}_{j,R}^{(i)}$ and $\hat{y}_{j,I}^{(i)}$ are the outputs of the LSTM/NN. In order to obtain the target values for the training of the LSTM/NN, we simulate 5×10^5 state discrimination measurements using Bayesian hypothesis testing with $\eta = 1.0$, $\xi = 0.997$, $\nu = 0$, and PNR(1). Thus, the training data consists of 5×10^5 experiments each represented by a sequence of detections $\{d_j\}_L$ and displacements $\{\beta_j\}_L$. The LSTM/NN can never *perfectly* implement the adaptive Bayesian strategy due to small differences between the target displacement values and the ones applied by the LSTM/NN. To this end,

we also add a small amount of Gaussian noise with standard deviation $\sigma = 0.01$ to the target values. This procedure makes the LSTM/NN more robust to deviations in the output β_j of the LSTM/NN compared to the actual Bayesian strategy, i.e. when $\beta_j \neq \alpha e^{i\phi_k}$.

We test the performance of the trained LSTM/NN by simulating 5×10^5 state discrimination experiments using the LSTM/NN to implement the measurement. After the L adaptive steps for a single measurement, the LSTM/NN updates a final time. The answer to the state discrimination problem θ_{disc} corresponds to the input phase ϕ_k that is closest to the phase of the final NN output $\arg(\beta_{L+1})$. For example, if the output of the NN after all updates is $(\mathrm{Re}[\beta_{L+1}] = -1.07, \mathrm{Im}[\beta_{L+1}] = 0.92)$, then the estimate for the input state will be the $|\alpha e^{i3\pi/4}\rangle$ even though $\arg(\beta_{L+1}) \approx 2.43$ and $3\pi/4 \approx 2.36$. We note that this procedure is just for finding θ_{disc} and that during the measurement the displacement is set to whatever values the NN outputs.

Figure 3.10 shows the simulation results when using the LSTM/NN to implement

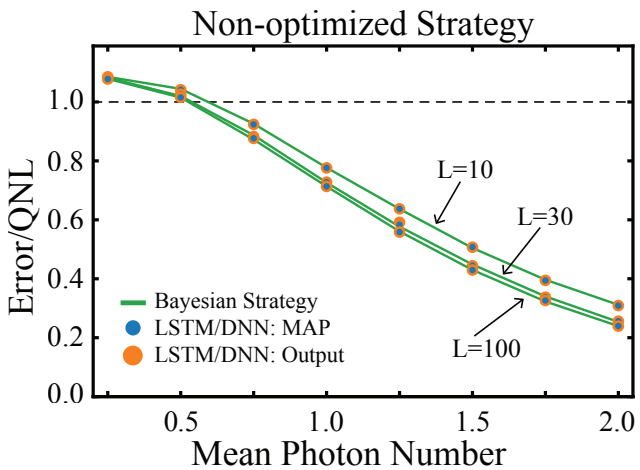

Figure 3.10: Results for implementing a non-optimized strategy with L=10, 30, and 100 adaptive steps which attempts to displace the input state to the vacuum state. The LSTM/NN is trained with the actual displacement values which implement the adaptive Bayesian strategy. The blue points show the error for the LSTM/NN strategy but with a maximum *a posteriori* (MAP) estimator for the final hypothesis and the orange points show the results for the LSTM/NN strategy but using the last displacement output to obtain the final hypothesis.

the state discrimination strategy for different input mean photon numbers $|\alpha|^2$ for L =10, 30, and 100 adaptive steps. The green line shows the results for the full Bayesian strategy which utilizes the photon counting likelihood functions. The blue points show the results for using the LSTM/NN to implement the measurement but using the maximum *a posteriori* (MAP) estimator to obtain θ_{disc}. The orange points show the result of using the LSTM/NN to implement the strategy and using the output of the NN to obtain θ_{disc} as described above, i.e. without any knowledge of the likelihood functions. We find that in this low mean photon number regime, the LSTM/NN is able to replace the full Bayesian strategy with no degradation in terms of the overall error probability. This result shows that the LSTM/NN has the potential to implement state discrimination measurements and perform equivalently to the Bayesian hypothesis testing approach.

3.3.3 Finding Optimized Strategies

We also aim to find optimized LSTM/NN-based strategies which can potentially surpass the QNL in the single-photon limit, as well as in parameter regimes (such as a large value of L) where optimizations of conventional Bayesian-based strategies are computationally impractical or even impossible. Furthermore, optimization is ideally performed solely based on sampling the performance of the LSTM/NN receiver, and adjusting the parameters of the network to achieve the best error probability. The entire LSTM/NN architecture can be represented by what we refer to as a solution vector V, which is comprised of all the parameters of the LSTM ($\mathbf{U}^{(l)}, \mathbf{W}^{(l)}, b^{(l)}$) and the NN weights and biases [142,143]. Thus, one would expect that there is a particular solution vector V_{opt} which allows the measurement to achieve the smallest error probability. To find V_{opt} without using the likelihood functions, we need to implement a so-called black-box optimization of the LSTM/NN [147]. This optimization strategy relies on repeatedly sampling the error probability by evaluating the LSTM/NN with a particular solution vector, and constructing new solutions based on the sampled error probability. Due to fact that any sampling will be inherently noisy, a well suited

algorithm is a genetic optimization [148] strategy known as differential evolution [149, 150]. This optimization algorithm, described in detail in Appendix C, relies on maintaining a pool of candidate solution vectors $\{V_i\}$ and updating each V_i depending on its corresponding error probability $P_E^{(i)}$ every iteration. If a candidate solution vector has a large error probability, then a replacement solution vector is constructed using combinations of "good" solution vectors, which have small error probabilities. In this way, the parameter space is explored through sampling the error probability and constructing new solution vectors based on the performance of each candidate. In addition, we use the machine learning technique of transfer learning [151] to obtain an initial value $V_i^{(0)}$ for each solution vector. Transfer learning is a technique where a network is trained to perform a task which is different but related to the original task, and then used as an initialization point [151,152]. Here, we use transfer learning by first training the LSTM/NN to implement a non-adaptive phase estimation measurement with a limited number of adaptive steps ($L = 5$). We then apply different instances of random noise to the "trained" network parameters to obtain the initial solution vector $V_i^{(0)}$ for each candidate. This process allows the LSTM/NN to learn a basic Bayesian-like property which can then be applied to the new problem of optimized state discrimination with a larger number of adaptive steps (see App. C for details).

Figure 3.11 shows the results for the differential evolution of the LSTM/NN as a function of the iteration number when using transfer learning for initialization of each solution vector. The green and black dashed lines show the QNL and Helstrom bound, respectively, and the red dashed line shows the error for a non-optimized strategy with $L = 10$, which attempts to displace the input state to the vacuum state. The blue dashed line shows the error for the ideal "historical" optimization strategy for $L = 10$ discussed in Section 3.2. The red solid line shows the error probability as a function of iteration number for the LSTM/NN with $L = 30$ adaptive steps. We find that the genetic algorithm can identify strategies which surpass the QNL after $\approx$25 iterations and the $L = 10$ non-optimized strategy after $\approx$50 iterations. The LSTM/NN can reach equivalent performance to the $L = 10$ globally optimized

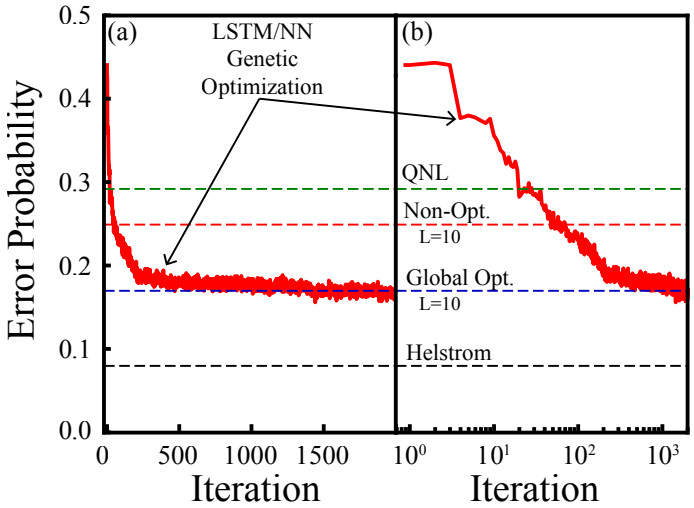

Figure 3.11: Results for the genetic optimization of the LSTM/NN architecture implementing a discrimination strategy with $L = 30$ at $|\alpha|^2 = 1.0$. Panel (a) and (b) show the simulation results on a linear and log scale in the horizontal axis for the optimization iteration. The LSTM/NN receiver (red solid) surpasses the QNL (green dashed) after only ≈ 30 generations, a non-optimized $L = 10$ strategy (red dashed) after ≈ 50 generations, and reach the error for a $L = 10$ "historical" optimization (blue dashed) after $\approx 10^3$ generations.

solution after ≈ 1500 iterations and then continues to slightly improve. This LSTM/NN receiver and genetic optimization can find optimized measurement strategies for large numbers of adaptive steps L. Furthermore, previous approaches based on "historical" optimizations discussed in Section 3.2 do not have favorable scaling as L increases. For example, a "historical" optimization strategy implementing $L = 10$ adaptive steps requires optimization of $2^{10} - 1 = 1023$ parameters, while a strategy with $L = 30$ adaptive steps requires optimization of over 10^9 parameters. However, the LSTM/NN is independent of L and the PNR of the detector and thus inherently scalable to more complex measurement strategies where "historical" optimization is impractical.

3.3.4 Future Directions in Reinforcement Learning

The main benefit of applying reinforcement learning to the LSTM/NN architecture is that knowledge of the noise properties of the communication channel is not required to find optimized measurement strategies. Thus, a direct application of this method

is for finding optimized measurement strategies when the communication channel has unknown noise associated with it. For example, if there is an unknown combination of phase diffusion and thermal noise, then reinforcement learning of the LSTM/NN receiver can find an approximately optimal measurement strategy. Furthermore, Chapter 2.2 discusses how the input alphabet for binary coherent states can be optimized jointly with the receiver to maximize information transfer in the presence of noisy communication channels [81]. In the context of noisy channels, this reinforcement learning framework can be extended to include optimization of the input alphabet for any number of states. This framework also allows for investigating the limits of information transfer of non-Gaussian channels where the fundamental limits are not well understood [13, 71, 135, 153, 154].

3.4 Conclusion

In this chapter, we presented different measurement strategies for discrimination of four coherent states of light. We fist investigated the capabilities of a "minimum" resource receiver for QPSK states [82] based on the proposal in Ref. [119]. These strategies are based on testing multiple hypotheses simultaneously using optimized displacements and PNR(1) detection. The strategy uses the smallest number of individual measurements, no photon number resolution, no complex feedback operations, and no phase stabilization, which makes it ideal for near-term high bandwidth communication. By using state-of-the-art optical elements and detectors [136], these measurements will be able to surpass the ideal QNL_{QPSK} and facilitate high bandwidth classical and quantum communication. Going beyond the "simple" measurement strategies, in Sec. 3.2 we implemented an adaptive receiver for QPSK states which can be globally optimized to achieve to best performance [83]. More complex optimizations can obtain better error probabilities and surpass the QNL_{QPSK} in the single photon regime where previous non-optimized demonstrations have fallen short. These optimized measurements are particularly

useful for quantum communication where the input powers are extremely small, and the receiver needs to be able to extract the most information possible. As an extension to these optimized receivers, in Sec. 3.3 we developed a reinforcement learning framework for optimized receivers for QPSK states where the Bayesian hypothesis testing strategy is replaced by a LSTM/NN architecture. The receiver strategy based on the LSTM/NN can approximate a non-optimized receiver strategy with no degradation in performance. Moreover, this LSTM/NN strategy can be optimized using reinforcement learning, which solely relies on sampling of the error probability instead of requiring access to the photon counting likelihood functions. The ability to construct an optimized strategy without specific knowledge of the likelihood functions and the communication channel will allow for finding optimized measurements which surpass the QNL, even in the presence of complex noise sources, and allow for investigating the attainable mutual information for non-trivial channels.

Chapter 4

Noise Tracking for Non-Gaussian Receivers

A receiver being able to efficiently identify and correct for noise in a communication channel is a critical barrier to reliably transferring information [50, 58, 99, 116]. In realistic communication channels, multiple static and dynamic noise sources may be present such as phase noise [58], amplitude noise, phase diffusion [81, 103, 104], or thermal noise [94, 155]. Ideally, channel noise is parameterized, estimated, and tracked by the receiver such that the damaging effects of the noise are counteracted. One natural solution to the problem of parameter tracking is for the sender to simply transmit a strong reference field in addition to the encoded information [156–160]. However, reference fields are not an efficient use of energy if the sender transmits high power optical pulses without taking that energy into account. To that end, conventional heterodyne receivers utilize a myriad of estimation and tracking algorithms to efficiently monitor and correct for channel noise with advanced digital signal processing without the need for reference pulses [58, 161–167]. However, these noise tracking algorithms usually rely on correcting heterodyne measurement data in post-processing [58, 168, 169].

Adaptive non-Gaussian receivers do not directly measure the real and imaginary

components of the input state, a property which conventional noise tracking methods with -dyne measurements rely on. Thus, conventional noise tracking methods cannot be applied to communication strategies using non-Gaussian receivers, which achieve sub-QNL sensitivities. As a result, non-Gaussian measurement strategies require fundamentally different approaches to parameter tracking which need to perform real-time estimation and correction, where the receiver strategy self-corrects to compensate for noise, as opposed to post-processing methods. In addition, parameter tracking strategies for non-Gaussian receivers must rely solely on the state discrimination data itself and not on extra resources such as reference pulses. Novel channel noise tracking methods must be developed and demonstrated such that non-Gaussian measurements can maintain their sub-QNL performance and become a practical and feasible technology.

We first address the problem of phase tracking by developing an intuitive estimator of time-dependent phase noise [85] in Section 4.1. We experimentally implement this strategy and show that sub-QNL performance can be maintained in the presence of applied and intrinsic phase noise. As an extension to the problem of phase tracking, we develop and simulate an approach based on a neural network estimator which implements parameter tracking of two parameters, phase and amplitude, simultaneously [86]. This approach based on a neural network can be efficiently implemented in hardware such as an FPGA, and operate at high bandwidths, which is required for maintaining sub-QNL performance in a realistic communication scenario.

4.1 Experimental Phase Tracking Demonstration

A particularly damaging type of channel noise for phase encoded coherent states is time-dependent phase noise due to phase changes between the sender and receiver, which degrades the performance of the receiver [58,68]. Phase noise in realistic channels is often modeled as a Gaussian random noise process [162, 164–167], which originates mainly from random phase differences between the transmitter laser and the LO

used by the receiver [50], as well as other sources [170–174]. Adaptive non-Gaussian measurements can achieve sub-QNL error rates [51, 117], but are highly sensitive to phase noise and must be actively corrected in real-time. Furthermore, conventional approaches to tracking and compensating for phase noise cannot be directly applied to adaptive non-Gaussian measurements. In this section, we develop and experimentally implement a phase tracking method which uses an estimator based on the photon counting statistics acquired from the non-Gaussian measurement. We show that this method can correct for Gaussian noise applied to the input states and allow for sub-QNL discrimination in the presence of moderate phase noise [85].

4.1.1 Phase Tracking Strategy

We use the adaptive photon counting strategy from Ref. [117] as a platform to study phase tracking in non-Gaussian receivers for QPSK coherent states $|\alpha_k\rangle = |\alpha e^{ik\pi/2}\rangle$ with $k \in \{0, 1, 2, 3\}$. Figure 4.1 shows a schematic of a non-Gaussian receiver operating below the QNL by implementing an adaptive, non-optimized strategy with $L = 7$ adaptive steps and PNR(3) [85, 117]. The receiver identifies any phase shift ϕ in the received states and feeds-forward a correction to the local oscillator (LO) by

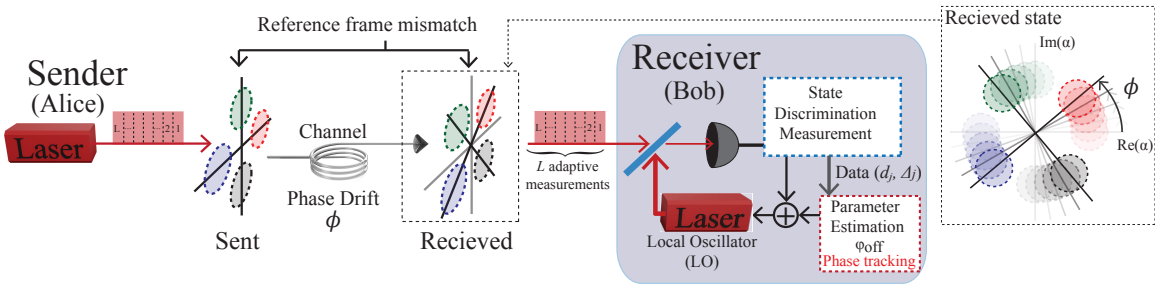

Figure 4.1: Phase tracking strategy where a receiver attempts to decode QPSK coherent states in the presence of phase noise. Unknown phase shifts ϕ acquired by the states when propagating through the communication channel will degrade the quality of the non-Gaussian receiver. Real-time parameter estimation is performed and a correction is fed-forward to the LO in order to identify and compensate for the phase noise. Figure adapted from: *Physical Review Research* 2, 023384 (2020).

implementing parameter estimation in real-time based on the state discrimination measurement data. Figure 4.2 shows the error probability of the strategy as a function of phase offset ϕ between the input states and LO for an input power of $|\alpha|^2$=2.0, 5.0, and 10.0 in blue, orange, and yellow, respectively. Strategies with higher powers increasingly outperform the QNL (gray lines) but are more sensitive to phase noise and the receiver can only tolerate $\phi \approx \pm 0.1$ rad. before the error becomes larger than the QNL for all three input powers. The receiver must perform active phase tracking in order to maintain sub-QNL error probabilities where estimates of the noise are fed-forward as a correction to the LO. Furthermore, noise tracking methods need to be computationally efficient enough to be experimentally implemented.

Efficient channel noise tracking must rely solely on the state discrimination measurement data, and not extra resources such as reference fields. To this end, the receiver constructs what we refer to as the detection matrix $\mathbf{D}$ to calculate an estimate $\hat{\phi}$ for the phase shift ϕ (or any noise in general) present in the photon counting measurement data. The size of the detection matrix is $M \times (m+1)$ where $M = 4$ for QPSK states and m is the photon number resolution. The receiver estimates the

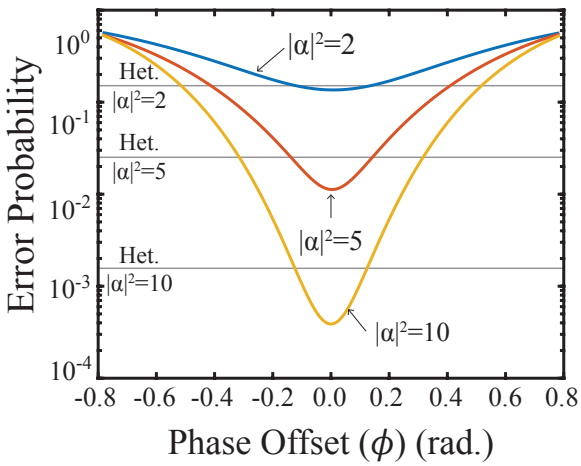

Figure 4.2: Error probability as a function of phase offset ϕ for strategies with an input mean photon number $|\alpha|^2$=2.0, 5.0, and 10.0 in blue, orange, and yellow respectively. Here the experimental imperfections are included in simulation where $\eta = 0.7$ and $\xi = 0.997$ with PNR(3). All the strategies can only tolerate $\phi \approx \pm 0.1$ rad. before the sub-QNL performance is lost. Figure adapted from: *Physical Review Research* 2, 023384 (2020).

phase noise using $\mathbf{D}$ after the completion of N state discrimination measurements. During each adaptive step j in a single state discrimination measurement, the phase of the LO $\arg(\beta_j) = \theta_j$ is changed between $\theta_j \in \{0, \pi/2, \pi, 3\pi/2\}$ to implement a recursive Bayesian strategy based on hypothesis testing [51, 85, 117]. The receiver provides an answer to the state discrimination problem θ_{disc} which corresponds to the most likely phase for the input state after the completion of L adaptive steps. The measurement also registers L photon detections $\{d_j\}_L$ in addition to the LO phases $\{\theta_j\}_L$. Due to the low error rate of the receiver, θ_{disc} is equal to the phase of the input state with high probability and the receiver uses θ_{disc} to infer the *relative* phase between the LO and the input at each adaptive step as: $\Delta_j = \theta_j - \theta_{disc}$. The data from the state discrimination measurement $\{\Delta_j, d_j\}$ is then binned into the detection matrix $\mathbf{D}$, which is initialized with all zeros. Each matrix element $\mathrm{D}_{k,l}$ is incremented at the end of every measurement by the total number of times the number of detected photons $d_j = l$ with the relative phase $\Delta_j = 2\pi k/M$ where $k \in \{0, 1, ..., M-1\}$. Thus, the rows of the detection matrix represent the sampled photon number distribution $P(n|\Delta_j = k\pi/2)$ for when the relative phase between the input state and LO is $k\pi/2$ for QPSK states. At the end of N state discrimination measurements, $\mathbf{D}$ contains the data from $L \times N$ samples of $\{\Delta_j, d_j\}$. The receiver then calculates an estimate $\hat{\phi}$ for the phase noise present in $\mathbf{D}$, feeds-forward a correction to the LO, and "resets" the detection matrix to zeros.

4.1.2 Sin-Cos (SC) Estimator

The receiver must utilize a particular estimator to identify the phase noise present in the measurement data. The input to the estimator is the matrix $\mathbf{D}$ and the output is an estimate $\hat{\phi}$ for the phase noise ϕ. Figure 4.3 shows a flowchart for the algorithm of the proposed Sin-Cos (SC) estimator which calculates $\hat{\phi}$ given $\mathbf{D}$. The blue box labeled "State Disc." shows how $\mathbf{D}$ is interpreted as samples of the photon number distributions for different relative phases Δ_j between the LO θ_j and the answer for the input state θ_{disc}. The orange box labeled "Phase Estimation" shows how the

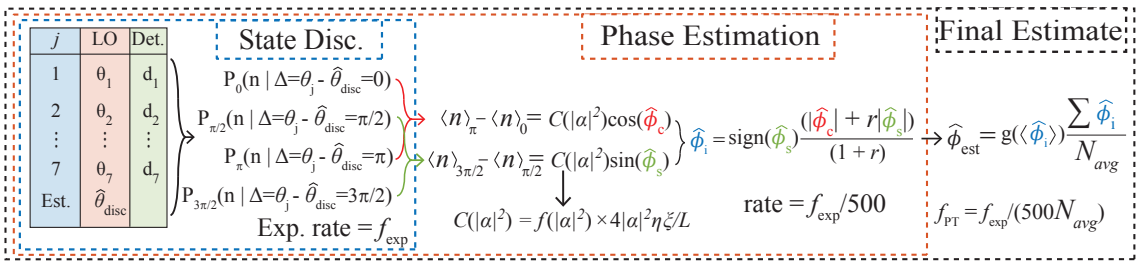

Figure 4.3: Flowchart of the proposed SC estimator described in the text. The data (blue box) from the previous N state discrimination measurements is used to construct raw estimates for the phase noise (orange box). Due to the large uncertainty, N_{avg} raw estimates $\hat{\phi}_i$ are averaged to obtain the final phase estimate $\hat{\phi}$ (black box). Figure adapted from: *Physical Review Research* 2, 023384 (2020).

algorithm calculates a raw estimate $\hat{\phi}_i$ after N state discrimination measurements. The SC estimator first takes the average across each row of $\mathbf{D}$, which is equivalent to finding the average value of each photon number distribution $P(n_j|\Delta_j = k\pi/2)$. Averaging across each row of $\mathbf{D}$ yields four mean values of $\langle n\rangle_0$, $\langle n\rangle_{\pi/2}$, $\langle n\rangle_\pi$, and $\langle n\rangle_{3\pi/2}$ for when $\Delta_j = 0$, $\pi/2$, π, and $3\pi/2$, respectively, which are ideally given by:

$$\langle n\rangle_0 = 2\eta\frac{|\alpha|^2}{L}\Big(1 - \xi\cos(\phi)\Big) + \nu, \tag{4.1}$$

$$\langle n\rangle_{\pi/2} = 2\eta\frac{|\alpha|^2}{L}\Big(1 - \xi\sin(\phi)\Big) + \nu, \tag{4.2}$$

$$\langle n\rangle_\pi = 2\eta\frac{|\alpha|^2}{L}\Big(1 + \xi\cos(\phi)\Big) + \nu, \tag{4.3}$$

$$\langle n\rangle_{3\pi/2} = 2\eta\frac{|\alpha|^2}{L}\Big(1 + \xi\sin(\phi)\Big) + \nu, \tag{4.4}$$

where η is the detection efficiency, ξ is the interference visibility of the displacement operation, and ν is the dark count rate per adaptive step. The algorithm obtains a value which is a sample of $\cos(\phi)$ by combining Equations (4.1) and (4.3) and in a similar way, obtains a sample of $\sin(\phi)$ from Equations (4.4) and (4.2). The SC estimator uses these samples as estimates for the phase noise as:

$$\langle n\rangle_\pi - \langle n\rangle_0 = C(|\alpha|^2)\cos(\phi) \rightarrow \hat{\phi}_c = \arccos\left[\frac{\langle n\rangle_\pi - \langle n\rangle_0}{C(|\alpha|^2)}\right], \tag{4.5}$$

$$\langle n\rangle_{3\pi/2} - \langle n\rangle_{\pi/2} = C(|\alpha|^2)\sin(\phi) \rightarrow \hat{\phi}_s = \arcsin\left[\frac{\langle n\rangle_{3\pi/2} - \langle n\rangle_{\pi/2}}{C(|\alpha|^2)}\right], \tag{4.6}$$

where $C(|\alpha|^2) = f(|\alpha|^2) \times 4|\alpha|^2\eta\xi/L$ and $f(|\alpha|^2)$ is a tunable factor which allows the estimator to account for the non-zero probability of error in discrimination [85]. The algorithm calculates a single raw estimate $\hat{\phi}_i$ from a weighted average of the two initial estimates $\hat{\phi}_c$ and $\hat{\phi}_s$:

$$\hat{\phi}_i = \text{sign}(\hat{\phi}_s)\frac{|\hat{\phi}_c| + r(|\alpha|^2)|\hat{\phi}_s|}{1 + r(|\alpha|^2)}, \tag{4.7}$$

where $r(|\alpha|^2)$ is the weighting factor which is used to maximize the linearity of $\hat{\phi}_i$ with respect to ϕ [85]. We set $N = 500$ in our demonstration such that a single raw estimate $\hat{\phi}_i$ contains the data from 500 state discrimination measurements.

The black box labeled "Final Estimate" shows that the last step in creating the final phase estimate $\hat{\phi}$ is to take the average of N_{avg} raw estimates $\hat{\phi}_i$ with an added gain factor g:

$$\hat{\phi} = g\left(\frac{1}{N_{avg}}\sum\hat{\phi}_i\right)\frac{1}{N_{avg}}\sum\hat{\phi}_i = g\big(\langle\hat{\phi}_i\rangle\big)\langle\hat{\phi}_i\rangle. \tag{4.8}$$

Ideally, the phase estimate $\hat{\phi}$ is linear as a function of the phase shift ϕ and has minimal uncertainty. However, in practice this is not the case and the values of $r(|\alpha|^2)$, $f(|\alpha|^2)$, and $g(\langle\hat{\phi}_i\rangle)$ are empirically tuned to best achieve the best performance for a particular input mean photon number $|\alpha|^2$. Reference [85] describes the detailed process for obtaining the optimal values of $r(|\alpha|^2)$, $f(|\alpha|^2)$, and $g(\langle\hat{\phi}_i\rangle)$ using Monte Carlo simulations. First, the optimal value of $f(|\alpha|^2)$ is one that minimizes the effect of errors in the state discrimination itself. Given that $P_E \neq 0$, there are errors in populating the sampled distributions $P(n|\Delta_j = k\pi/2)$ (rows of $\mathbf{D}$) and therefore their mean values $\langle n\rangle_{k\pi/2}$. Discrimination errors will mainly cause $\langle n\rangle_\pi - \langle n\rangle_0 < 4|\alpha|^2\eta\xi/L$, which biases $\hat{\phi}_c$ away from zero when there is no phase noise. These errors also cause a bias in $\hat{\phi}_s$ but *towards* zero and we tune the value of $f(|\alpha|^2)$ to minimize the effect of the errors. The optimal weighting factor $r(|\alpha|^2)$ minimizes the difference between ϕ and $\hat{\phi}_i$ at the endpoints of the capture range R of the phase tracking, where $R = \pm 0.6$ rad. in our experiment. The limited range R is due to requiring a relatively small resolution of ≈ 5 mrad. when discretizing the LO phase correction into 256 phase

points (1.2 rad./256≈ 5 mrad.). The optimal gain function $g(\langle\hat{\phi}_i\rangle)$ maximizes the linearity of $\hat{\phi}_i$ as a function of ϕ across the capture range R once $f(|\alpha|^2)$ and $r(|\alpha|^2)$ are determined and is found empirically with Monte Carlo simulations.

The performance of the SC estimator can now be examined after finding optimal values for each of the tunable parameters in the SC estimator for different input mean photon numbers. We benchmark the SC estimator against the performance of a Bayesian estimator, which is calculated using the complete photon counting likelihood functions as:

$$\hat{\phi}_B = \arg\left(\int_{-\pi}^{\pi} e^{i\phi} P(\phi|\mathbf{D}) d\phi\right), \tag{4.9}$$

where $P(\phi|\mathbf{D})$ is the posterior distribution for the phase ϕ given the state discrimination measurement data $\mathbf{D}$ and is obtained by:

$$P(\phi|\mathbf{D}) = \frac{P(\mathbf{D}|\phi)P(\phi)}{P(\mathbf{D})} = \mathcal{N}P(\phi)\prod_{k=0}^{M-1}\prod_{l=0}^{m}\mathcal{L}(l|\phi - k\pi/2)^{D_{k,l}}, \tag{4.10}$$

$$\mathcal{L}(l|\psi) = \frac{1}{l!}\langle n(\psi)\rangle^l e^{-\langle n(\psi)\rangle}, \tag{4.11}$$

$$\langle n(\psi)\rangle = 2\eta|\alpha|^2\Big(1 - \xi\cos(\psi)\Big) + \nu, \tag{4.12}$$

where the total likelihood function $P(\mathbf{D}|\phi)$ is the product of each individual photon counting likelihood function $\mathcal{L}(l|\phi - k\pi/2)$ for each matrix element $D_{k,l}$.

Figure 4.4(a-c) shows the uncorrected ($g = 1$) SC estimator (orange) and the corrected ($g \neq 1$) SC estimator (green) compared to an uncorrected (blue) and corrected (black) Bayesian estimator for $|\alpha|^2$=2.0, 5.0, and 10.0, respectively. Solid lines represent the average across 100 different Monte Carlo samples and the shaded regions correspond to one standard deviation. The uncorrected estimates are biased towards zero as the phase offset increases for all three input powers. The purpose of the gain function $g(\cdot)$ is to "undo" this bias such that on average $\hat{\phi} = \phi$ across the entire capture range $R = \pm 0.6$ rad. The corrections cannot overcome the effect of a large P_E for $|\alpha|^2 = 2.0$ beyond $\phi \approx \pm 0.5$ rad., as can be seen by the corrected estimates for both estimators deviating from being linear.

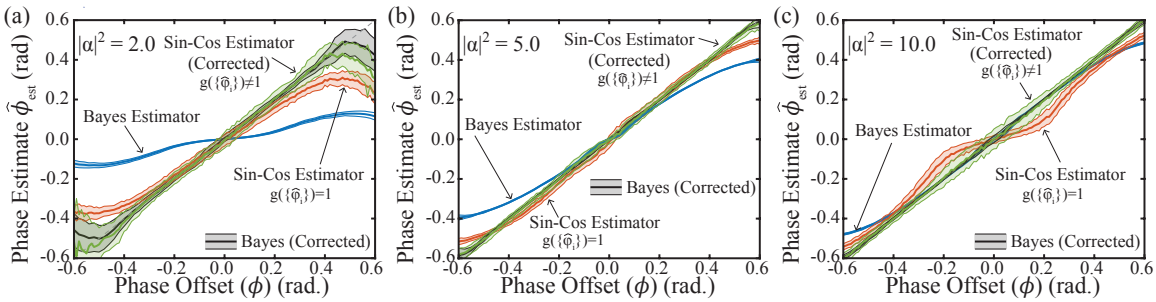

Figure 4.4: Linearity of the SC estimator for different applied phase offsets ϕ for $|\alpha|^2 = 2.0$ (a), 5.0 (b), and 10.0 (c). A Bayesian estimator (blue & black) and the SC estimator (orange & green) can both be corrected such that they are linear with respect to the actual phase offset. This correction is done by applying a custom gain function g which is found empirically through Monte Carlo simulations. The solid lines represent the average of 100 Monte Carlo samples and the shaded regions correspond to one standard deviation. Figure adapted from: *Physical Review Research* 2, 023384 (2020).

4.1.3 Experimental Implementation

Figure 4.5 shows the experimental setup used to implement phase tracking with the SC estimator. The phase of the signal and LO are each controlled by a phase modulator (PM) and after interference on a 99/1 beam splitter, the displaced input state $\hat{D}(\beta_j)|\alpha_k\rangle$ is detected by a single photon detector (SPD). We use an Altera Cyclone II field programmable gate array (FPGA1) to implement the state discrimination strategy with $L = 7$ adaptive steps and PNR(3) at a repetition rate of $f_{exp} \approx 12$ kHz [85, 117].

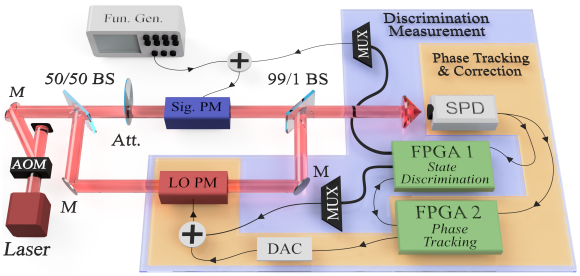

Figure 4.5: Experimental setup for phase tracking in adaptive non-Gaussian receivers. The phase noise is applied to the signal phase modulator (PM) using a function generator and two field programmable gate arrays (FPGAs) implement the state discrimination measurement and SC estimator. The phase correction is fed-forward to the LO PM. Figure adapted from: *Physical Review Research* 2, 023384 (2020).

We achieve a detection efficiency of $\eta = 0.72$ and an interference visibility of $\xi = 0.998$ for this experiment. A second FPGA (FPGA2) contains the estimation algorithm described above to implement the SC estimator. FPGA2 receives the same state discrimination data and calculates and applies the phase correction. An 8-bit digital-to-analog converter (DAC) controls the LO phase correction such that there are $2^8 = 256$ discrete values which can be applied across the capture range $R = \pm 0.6$ rad. A function generator connected to the phase modulator in the signal arm applies phase noise to the phase of the input states. In addition, the relative phase between the input and LO can be "unlocked" such that the total phase noise is a combination of the applied noise and the natural drift of the experimental setup.

4.1.4 Experimental Results

We investigate the capability of phase tracking method based on the SC estimator by first applying fixed phase shifts ϕ_{app} to the phase modulator in the signal arm of the experiment. We then study the dynamic tracking ability of the SC estimator by applying Gaussian phase noise to the signal. We apply different strengths and bandwidths of noise for different input powers and also study the performance of the estimator as a function of N_{avg}, the number of averages used for a single estimate.

Constant Phase Offset

We first study the ability of the phase tracking algorithm to identify and correct sudden, fixed phase shifts, or phase "jumps". We apply a constant phase offset ϕ_{app} to the input states in these experiments and then, after a fixed amount of time, the receiver enables the phase tracking algorithm such that the offset is corrected. Figure 4.6(a) shows the probability of error as a function of time as the phase jump is applied and corrected. The input mean photon number is $|\alpha|^2 = 5.0$ and we set the number of averages per estimation to $N_{avg} = 20$, giving a phase tracking frequency of $f_{PT} = f_{exp}/(500 \times 20) \approx 1.2$ Hz. The solid lines and shaded regions correspond to the

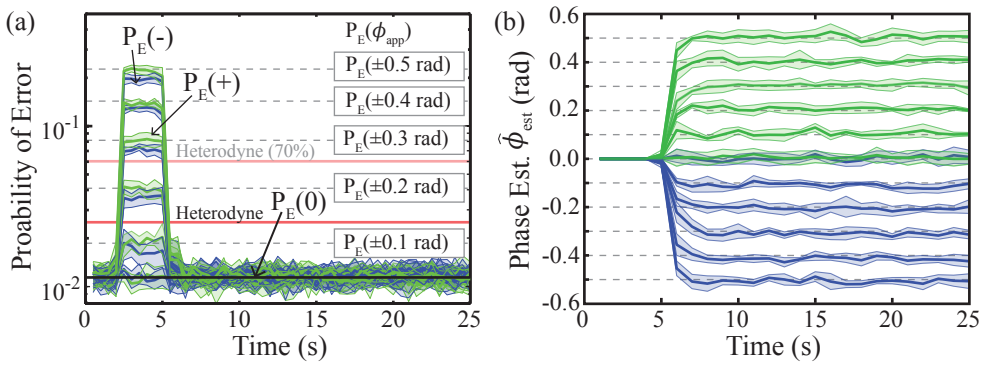

Figure 4.6: (a) Experimental performance of the SC estimator based phase tracking method when applying fixed phase shifts ϕ_{app} to the input states. The solid lines and shaded regions correspond to the average and standard deviation of the error probability across five independent runs of the experiment. Initially, no phase noise is applied and the phase tracking algorithm is disabled. At $t = 2$ s, a phase shift ϕ_{app} is applied and causes the error probability to increase depending on $|\phi_{app}|$. At $t = 5$ s, the SC algorithm is enabled and can identify and correct the phase shift in only one or two iterations. (b) Estimated phases $\hat{\phi}$ (colored lines) compared to the applied phase shifts ϕ_{app} (dashed lines). Figure adapted from: *Physical Review Research* 2, 023384 (2020).

average and standard deviation of the error probability across five independent runs of the experiment. Starting at time $t = 0$ s, the error probability is calculated across time bins of approximately 0.5 s each, corresponding to 5×10^3 ($\approx f_{exp} \times 0.5$ s) experiments. Initially, there is no applied phase offset ϕ_{app} and the phase tracking algorithm is turned off in order to verify that the measurement is performing as expected with error probability $P_E(\phi_{app} = 0)$ (solid black), below the QNL (solid red). The phase jumps are applied at time $t = 2$ s with different magnitudes and signs from $\phi_{app} = \pm 0.1$ to ± 0.5 rad. (dashed gray). The green (blue) lines show the average error probability for when the positive (negative) phase shifts are applied. The receiver enables the phase tracking algorithm at time $t = 5$ s and within two estimation iterations, the algorithm corrects for the phase offset accurately enough for the receiver to maintain the same error probability as with no phase noise $P_E(\phi_{app} = 0)$. Figure 4.6(b) shows the estimated phase shifts (green and blue solid lines) as a function of time compared to the applied phase shifts ϕ_{app} (gray dashed lines). The phase tracking algorithm can efficiently identify and correct for the applied phase noise and maintain performance below the QNL.

Phase Tracking at Different Noise Strengths & Frequencies

To fully demonstrate the capabilities of the proposed algorithm, we implement real-time phase tracking for different input powers, noise strengths, and noise bandwidths. We note that for these experiments, the interferometer was purposefully unlocked in addition to the applied phase noise to represent a realistic situation [58, 166, 167]. We simulate discrete Gaussian random walks as a phase noise model with strength σ_1 and frequency $f_{RW} = 100$ Hz ($\Delta T_{RW} = 10$ ms). We apply the random walks to the phase of the input states as: $\phi_{app}(\tau + \Delta T_{RW}) = \phi_{app}(\tau) + \mathcal{N}(0, \sigma_1^2)$ where τ is time discretized into steps of ΔT_{RW}.

Figure 4.7(a-c) shows the probability of error for a strategy with $|\alpha|^2 = 5.0$ and $N_{avg} = 20$ when applying Gaussian phase noise of different strengths. We tune the

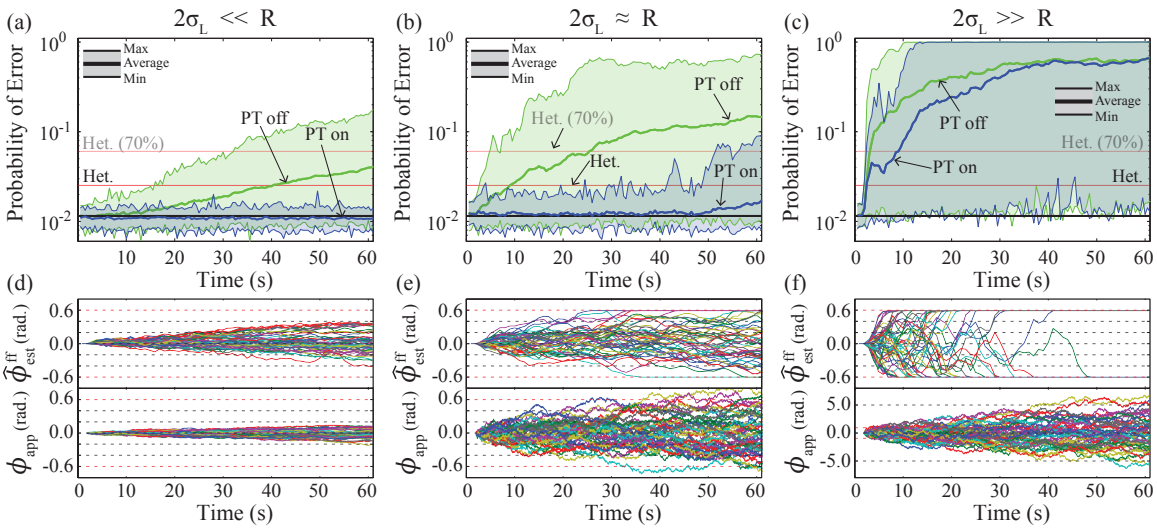

Figure 4.7: Experimental performance of the phase tracking algorithm as a function of time for $|\alpha|^2 = 5.0$ and $N_{avg} = 20$ when applying Gaussian random walks in phase to the input states. The strength of the phase noise varies from (a) small ($2\sigma_T \ll R$), to (b) moderate ($2\sigma_T \approx R$), to (c) severe ($2\sigma_T \gg R$). The sub-QNL performance is maintained in the presence of moderate phase noise until the estimated waveforms reach the phase tracking range $R = \pm 0.6$ rad., where the feed-forward corrections are locked to R. The effect of clamping to the range R is an increase in the error probability starting at $t \approx 50$ s, but the increase is solely due to the limited value of R in our implementation. Figure adapted from: *Physical Review Research* 2, 023384 (2020).

noise strength σ_1 such that the total deviation of the random walks $\sigma_T = \sqrt{N_T}\sigma_1$ after $N_T = 6500$ time steps ΔT_{RW} corresponds to (a) $\sigma_T \ll R/2$ ($\sigma_1 = 0.1$ mrad.), (b) $\sigma_T \approx R/2$ ($\sigma_1 = 5$ mrad.), and (c) $\sigma_T \gg R/2$ ($\sigma_1 = 25$ mrad.). The blue and green solid lines show the error with (PT on) and without (PT off) the phase tracking algorithm enabled, respectively, and the shaded regions correspond to the maximum and minimum errors due to the phase noise. The lower panels (d-f) show the corresponding applied phase noise $\phi_{app}(\tau)$ and phase corrections ϕ^{ff} which are fed-forward to the LO phase.

Figure 4.7(a) shows the uncorrected error probability slowly drifting upwards and the phase tracking algorithm can easily maintain the sub-QNL error rate when enabled for the duration of the phase walks (≈ 65 s). By comparing the applied and estimated phase noise in (d), the spread of ϕ^{ff} is much greater than the applied noise $\phi_{app}(\tau)$. These deviations are due to the fact that the experiment is unlocked and "drifting" such that natural phase drifts in the setup are also tracked. Figure 4.7(b) shows that moderate strength phase noise quickly degrades an uncorrected experiment. Enabling the PT algorithm allows for discrimination below the ideal QNL and error rates almost equal to the expected error in the absence of phase noise. After $t \approx 50$ s, the error probability starts to increase as well as the spread of errors (shaded region). This increase in the error is solely due to the finite capture range $R = \pm 0.6$ rad. of our implementation of the phase tracking algorithm. Once the estimated waveform goes beyond ± 0.6 rad., the applied correction ϕ^{ff} which is fed-forward to the LO is clamped to ± 0.6 rad. and the effect can be seen in the panels (e) and (f). Figure 4.7(c) shows that severe phase noise immediately degrades the uncorrected measurement and the corrected measurement appears to track the noise briefly before encountering the capture range R. By examining the performance before encountering the capture range, it appears that even with severe noise, the receiver potentially maintains an error rate in between the ideal and equivalent QNL.

We also investigate the capabilities of the SC estimator based algorithm for higher and lower input powers in addition to $|\alpha|^2 = 5.0$. Figure 4.8(a) shows the error

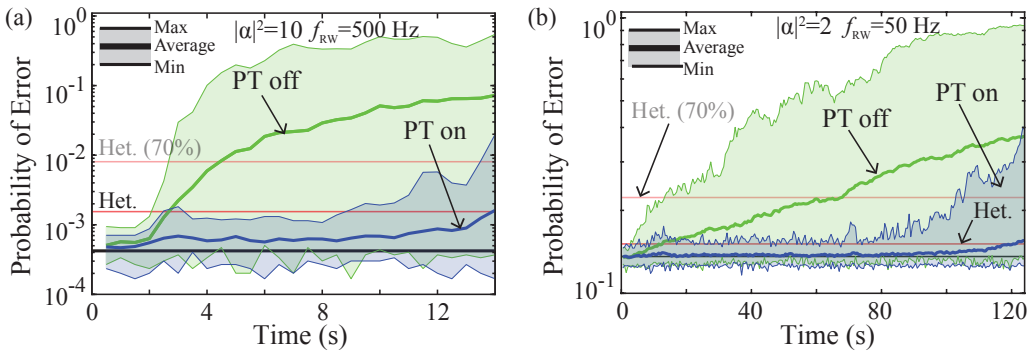

Figure 4.8: Experimental performance of the phase tracking algorithm for different input mean photon numbers of (a) $|\alpha|^2 = 2.0$ and (b) $|\alpha|^2 = 10.0$. Due to the different error probabilities for the different input powers, the number of averages needs to be changed to $N_{avg} = 40$ (4) when $|\alpha|^2$=2.0 (10.0). The phase tracking bandwidth is also changed accordingly to f_{RW}=50 Hz for $|\alpha|^2 = 2.0$ and f_{RW}=500 Hz for $|\alpha|^2 = 10.0$. Figure adapted from: *Physical Review Research* 2, 023384 (2020).

probability for an input power of $|\alpha|^2 = 10.0$ in the presence of moderate phase noise. The SC estimator requires less averaging ($N_{avg} = 4$) to achieve acceptable uncertainties due to the lower probability of error for $|\alpha|^2 = 10.0$. This change in N_{avg} also means that the algorithm should be able to track higher frequency noise. To this end, we set $f_{RW} = 500$ Hz while leaving $\sigma_1 = 5$ mrad. such that $2\sigma_T \approx R$ after 13 s. We find the phase tracking algorithm enables sub-QNL state discrimination in the presence of phase noise which would otherwise cause the performance to be entirely degraded. Figure 4.8(b) shows the error probability for an input mean photon number of $|\alpha|^2 = 2.0$. We increase the number of averages to $N_{avg} = 40$ due to the relatively high but still sub-QNL error rate. We also decrease the frequency of the phase noise accordingly to $f_{RW} = 50$ Hz such that $2\sigma_1 \approx R$ after 130 s. Similar to other input powers, the SC estimator based algorithm is able to maintain superior performance of the adaptive non-Gaussian receiver in the presence of phase noise.

4.1.5 SC Estimator Performance for Different N_{avg}

A critical aspect of any phase tracking method is how the performance changes with the number of samples $N \times N_{avg}$ used for estimation. Increasing the number of

averages N_{avg} reduces the estimation uncertainty $(\Delta\hat{\phi})^2$, but at the cost of increasing the time needed to calculate a single estimate $\hat{\phi}$. Increasing N_{avg} (and subsequently the estimation time) will also increase the deviation of the random phase walks during the estimation time, leading to an increase in the error probability. On the other hand, reducing the value of N_{avg} to provide faster estimation reduces the deviation of the random phase walks, but also increases the intrinsic variance of the SC estimator such that $(\Delta\hat{\phi})^2$ dominates the overall estimation uncertainty. Thus, there exists an optimal number of averages N_{avg} required to calculate each estimate for a specific phase noise strength. This optimal N_{avg} is a value which balances the contributions to the overall uncertainty from the intrinsic variance and the noise variance.

Figure 4.9(a-c) shows examples of estimates $\hat{\phi}$ from the SC estimator for $|\alpha|^2 = 5.0$ when no phase noise is applied ($\phi = 0$) for N_{avg}=2, 15, and 40, respectively, as well as a histogram of the estimates on the right of each plot. The estimates have a large uncertainty for a small value of N_{avg}, and increasing N_{avg} correspondingly decreases the estimator variance. Figure 4.9(d) shows the variance σ_0^2 of the SC estimator (blue points) as a function of N_{avg} calculated from the data shown in (a-c), where the points and error bars show the average and standard deviation across 5 experimental data sets. The inset (i) shows the data on a log-log scale where a model of $\sigma_0^2 = a/N_{avg}$ fits the data well, with a fitted coefficient of $a = 0.014$. The black line shows the deviation variance σ_{RW}^2 of a Gaussian random walk process as a function of N_{avg}, which is linear in time.

Figure 4.9(e) shows an example of the performance of the algorithm when tracking a single applied random walk (black line) with σ_1=5 mrad. for when N_{avg}=2 (gray), N_{avg}=10 (orange), and N_{avg}=40 (blue). Figure 4.9(f) shows the difference $\Delta = \hat{\phi} - \phi_{RW}$ between the estimated phase noise and applied phase noise for the different values of N_{avg} from (e). The deviations from zero have contributions from the intrinsic uncertainty of the SC estimator as well as the random phase noise. The variance when $N_{avg} = 2$ (gray) is mostly due to the uncertainty of the SC estimator σ_0^2 and instead of the phase noise σ_{RW}^2. On the other hand, the deviations of the random walks dominates

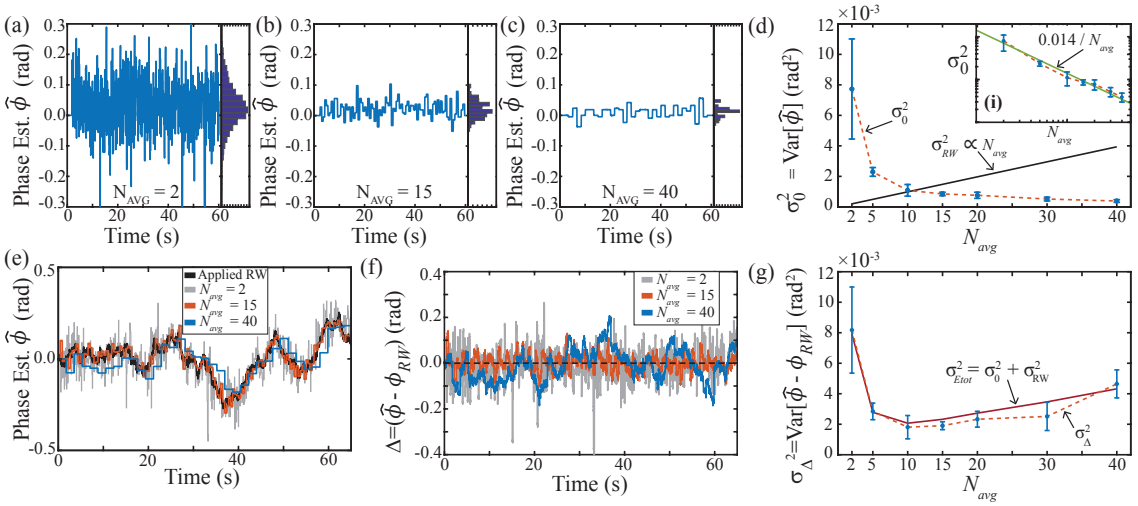

Figure 4.9: Study of the SC estimator performance as a function of N_{avg} for $|\alpha|^2$=5.0. (a-c) Distribution of estimates when $\phi = 0$ for N_{avg}=2, 15, and 40, respectively. (d) Variance of the estimates as a function of N_{avg}, which decreases as more samples are used for estimation. The inset shows the expected $1/N_{avg}$ statistical scaling of the estimator. (e) Estimates of a single applied random walk (black) for N_{avg}=2 (gray), N_{avg}=15 (orange), and N_{avg}=40 (blue), as a function of time. (f) Difference Δ between the estimated waveform and applied phase noise in (e). (g) Total estimator variance σ_Δ^2 from (f) which is a combination of the intrinsic estimator variance σ_0^2 from (d) and the deviation variance σ_{RW}^2 from the Gaussian random walks in phase, which scales linearly with N_{avg}. There is an optimal value $N_{avg} \approx 10$ which balances the two contributions to the overall variance. Figure adapted from: *Physical Review Research* 2, 023384 (2020).

the overall uncertainty when $N_{avg} = 40$ due to the relatively low phase tracking bandwidth compared to the phase noise. Figure 4.9(g) shows the total phase variance σ_Δ^2 (blue points) of Δ from (f) as a function of N_{avg} across 5 different noise realizations. The red solid line shows the expected total variance σ_{Etot}^2 which is equal to the sum of the SC estimator variance σ_0^2 plus the deviation variance of the random walk σ_{RW}^2. This expected variance has the form of $\sigma_{Etot}^2 = a/N_{avg} + bN_{avg} = a/N_{avg} + (500\sigma_1^2)N_{avg}$, and shows good agreement with the measured total variance. If the receiver knows the strength of the random phase noise σ_1^2, then an optimal value of N_{avg} is one which minimizes the total variance and is $N_{avg} \approx 10$ in (g). One would expect the optimal number of averages to be $N_{avg} = \sqrt{a/b}$ for a general level of noise b and estimator scaling a. Thus, if the estimator variance a/N_{avg} decreases or if the noise variance bN_{avg} increases, then the optimal N_{avg} will be smaller. On the other hand, if the noise

variance decreases or the estimator variance increase, then the optimal N_{avg} be need to be larger, thus reducing the achievable phase tracking bandwidth.

4.1.6 Discussion

Implementing efficient and accurate methods for phase tracking will allow adaptive non-Gaussian strategies to be transitioned out of the lab. Such techniques will enable error rates beyond the QNL in practical and realistic communication systems. On one hand, a Bayesian estimator will yield small estimation uncertainty but is very computationally expensive to calculate. On the other hand, simple estimators such as the SC estimator can be easily and efficiently calculated at the bandwidths necessary for realistic implementations. We studied the capabilities and limitations of the SC estimator across different phase noise regimes and input powers. In principle, the SC estimator can be extended to multi-parameter tracking but may not be able to achieve small estimation error at the required bandwidth. However, more advanced techniques may be leveraged to develop noise tracking methods for non-Gaussian receivers which implement high dimensional parameter tracking at high bandwidth.

4.2 Channel Noise Tracking with Neural Networks

In addition to phase noise, a communication channel can also exhibit amplitude noise such that the power of the received states fluctuates in time. Compensating for channel noise by actively correcting a non-Gaussian receiver now requires estimation and tracking of two parameters, a more complex problem where simple methods may not perform well. Recently, machine learning techniques have proven to be useful tools for solving problems in quantum information [175–183] and optical communication [184–187], where conventional methods may be computationally inefficient [188]. A particularly useful approach in optical communication is using artificial neural networks [189, 189–194] for channel noise estimation and monitoring [195–197]. Here,

we develop a method for tracking multiple noise parameters simultaneously based on a neural network estimator to correct adaptive non-Gaussian receivers [86]. We show using Monte Carlo simulations that this method can efficiently correct for phase and amplitude noise in parameter regimes where realistic implementations can potentially maintain sub-QNL performance.

4.2.1 Noise Tracking Strategy

Figure 4.10(a) shows the proposed method for parameter tracking based on a neural network (NN) estimator. The sender (Alice) prepares QPSK encoded coherent states and transmits the states across a channel to the receiver (Bob). However, there is dynamic phase and amplitude noise imparted onto the coherent states by the channel such that they arrive at the receiver with a time-dependent amplitude and phase. The receiver attempts to perform state discrimination of the QPSK states using an adaptive non-Gaussian measurement, but the channel noise severely degrades the achievable performance of the measurement. Estimating the phase and amplitude noise using solely the measurement data allows the receiver to correct the state discrimination measurement and maintain the sub-QNL performance.

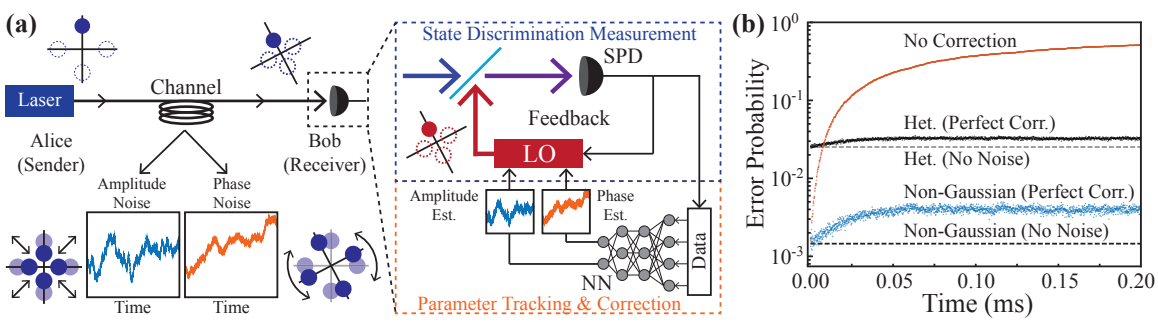

Figure 4.10: (a) NN-based multi-parameter noise tracking strategy where phase and amplitude noise is estimated and tracked by the adaptive non-Gaussian receiver. (b) Simulated error probability when applying phase and amplitude noise with no correction (orange), and perfect correction (blue), compared to a heterodyne receiver with perfect correction (black). The dashed black and gray lines show the error for non-Gaussian and heterodyne receivers in the absence of channel noise. Figure adapted from: *Physical Review Research* 3, 013200 (2021).

The receiver implements an adaptive measurement based on photon counting with $L = 10$ and PNR(10) following a hypothesis testing procedure where the phase of the displacement $\hat{D}(\beta_j)$ in adaptive step j is $\arg(\beta_j) = \theta_j \in \{0, \pi/2, \pi, 3\pi/2\}$ with magnitude $|\beta_j| = |\alpha|$ [51, 117]. We simulate a receiver with ideal detection efficiency ($\eta = 1.0$) and zero dark counts ($\nu = 0$), but a non-ideal interference visibility $\xi = 0.997$ to represent a realistic implementation but with an ideal detector. The relative phase $\phi(t)$ between the input and LO as well as the input energy $|\alpha|^2(t)$ are not constant in time due to channel noise. The receiver uses the measurement data and the answer to the discrimination problem θ_{disc} to implement phase and amplitude tracking, where estimates for the channel noise are fed-forward to correct the LO and maintain sub-QNL performance.

Figure 4.10(b) shows an example of the error probability of the receiver as a function of time for an average input state energy of $\langle \hat{n} \rangle_0 = |\alpha|^2(0) = 5.0$ averaged over 5000 realizations of channel noise. The orange points show the error rate when the receiver does not implement noise tracking and blue points show the error probability when the receiver has perfect knowledge of the noise, corresponding to the best possible performance. Perfect tracking refers to the receiver having complete knowledge of the time-dependent amplitude and phase of the input states such that the measurement can be corrected by setting $\arg(\beta) = \theta_j + \phi(t)$ and $|\beta|(t) = |\alpha|(t)$. The black points show the performance of a heterodyne receiver also with perfect correction, which serves as the effective QNL. The dashed lines show the error probability in the absence of noise for heterodyne (gray) and non-Gaussian (black) receivers. The error for the non-Gaussian receiver remaining below the effective QNL shows that the measurement maintains sub-QNL performance if the receiver implements effective channel noise tracking and correction.

The receiver collects the state discrimination measurement data in the detection matrix $\mathbf{D}$, which is a $M \times (m + 1)$ dimensional matrix. This matrix contains the data from N state discrimination measurements and is used for estimation of the noise imparted by the channel over the previous N measurements, as described in Sec. 4.1.1.

The discrimination of a single input state results in a sequence of L photon detections $\{d_j\}_L$ and relative phases $\{\Delta_j\}_L$ between the LO and input at each adaptive step. The receiver uses the final estimate θ_{disc} for the input state to infer $\Delta_j = \theta_j - \theta_{disc}$ and arranges the measurement data into $\mathbf{D}$ based on $\{d_j, \Delta_j\}$ such that the rows of $\mathbf{D}$ represent the sampled photon number distributions $P(n|\Delta_j = k\pi/2)$ with $k \in \{0, 1, 2, 3\}$ for QPSK states. The receiver must utilize a particular estimator in order to extract the noise information from $\mathbf{D}$ after N state discrimination measurements. A Bayesian estimator (see Sec. 4.1.2) which uses the full photon counting likelihood functions will have small uncertainty, but is computationally expensive to calculate. Furthermore, the Sin-Cos estimator described in Sec. 4.1.2 has large uncertainty and is not easily scalable to two (or more) dimensions for multi-parameter tracking.

4.2.2 Neural Network Architecture and Algorithm

We construct and train a neural network to act as a versatile estimator of the input phase and amplitude noise to solve the issue of computational cost while achieving small uncertainty. One benefit of utilizing a NN is that it is far more amenable to a practical implementation in devices such as FPGAs compared to a Bayesian estimator. Figure 4.11 shows the NN architecture which has ten layers (eight hidden) where each layer has a Leaky ReLU activation function [198] except the output layer, which has a sigmoid activation. The input to the NN is the detection matrix $\mathbf{D}$ first normalized across each row, and then arranged into a 1D vector ($D_{i,j} \rightarrow D_{i(m+1)+j}$) along with the LO intensity for the previous N measurements. We denote the time dependent input intensity of the QPSK states as $\mathcal{A}(\tau) = |\alpha|^2(\tau)$ and the LO intensity as $\mathcal{B}(\tau) = |\beta|^2(\tau)$. Here τ represents time discretized into steps of ΔT, where $1/\Delta T$ is the repetition rate, which we set $1/\Delta T = 100$ MHz ($\Delta T = 10$ ns) in all simulations to represent an achievable near-term communication bandwidth for integrated non-Gaussian receivers [199]. The outputs of the NN are raw estimates $\hat{\mathcal{A}}_{NN}$ and $\hat{\phi}_{NN}$ for the input intensity $\mathcal{A}(\tau)$ and phase offset $\phi(\tau)$ for the previous N state discrimination measurements, respectively.

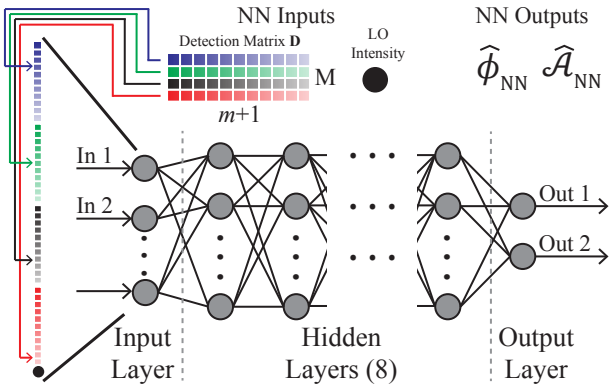

Network parameters	
Number of Layers	10 (8 Hidden)
Hidden layer size (l_i)	$\{32, 32, 32, 32, 16, 16, 8, 8\}$
Activation function	Leaky ReLU ($a = 0.1$)
Initialization	Norm(0, $\sigma^2 = 1/l_i$) (Xavier)
Training parameters	
Cost function	Weighted mean squared error
Optimizer	RMSprop (momentum=0.8)
Learning rate	50×10^{-6}
Epochs	2000

Figure 4.11: Neural network architecture where the state discrimination data (detection matrix $\mathbf{D}$) is fed into the network which outputs raw estimates ($\hat{\phi}_{NN}, \hat{\mathcal{A}}_{NN}$) for the phase $\phi(t)$ and amplitude $\mathcal{A}(t)$ noise. The network is made of ten layers (eight hidden) each with a Leaky ReLU activation except for the output later, which has a sigmoid activation. Figure adapted from: *Physical Review Research* 3, 013200 (2021).

The NN is trained using 5×10^5 training data points, which are obtained using Monte Carlo simulations. Each training data point consists of N state discrimination measurements where input intensity $\mathcal{A}$ and LO intensity $\mathcal{B}$ are constant for the N measurements and sampled from a uniform distribution $\mathcal{U}(0.05, 25)$. The phase offset for the N measurements is also constant and sampled from a zero-mean Gaussian distribution with variance $\sigma^2 = 0.25$. We also randomly sample the number of state discrimination measurements N used for each point from a uniform distribution $\mathcal{U}(2, 200)$. This random sampling allows a single NN to be used across a wide range of channel noise parameters as well as different values of N. The target values $y_{i,j}$ (i.e. what we want the NN to output) in the training are the actual input intensity ($y_{i,1}$) and phase offset ($y_{i,2}$), and we use the RMSprop optimizer [200] with a weighted mean-square-error cost function: $C = \sum_{i,j} w_i (y_{i,j} - \hat{y}_{i,j})^2$. The weight w_i of the i^{th} sample is given by $w_i = e^{-(\mathcal{A}_i - \mathcal{B}_i)^2/2}$ such that samples where the LO intensity is close to the input intensity are emphasized while building robustness to large amplitude fluctuations.

We use the NN as a two dimensional estimator of channel phase and amplitude noise. As a model for phase noise, we simulate a Gaussian random walk in phase.

A single step of a random walk is given by a sample from a zero-mean Gaussian distribution with variance $\sigma_1^2 = 2\pi\Delta\nu\Delta T$ such that:

$$\phi(\tau + \Delta T) = \phi(\tau) + \mathcal{N}(0, \sigma_1^2), \tag{4.13}$$

where $\Delta\nu$ is the phase noise bandwidth originating from laser phase noise [58, 162, 165], or other noise sources [172]. We simulate noise in the input intensity $\mathcal{A}(\tau)$ and use an Ornstein-Uhlenbeck (OU) process as a model for amplitude noise [201, 202], which is described by the stochastic differential equation:

$$\mathcal{A}(\tau + \Delta T) - \mathcal{A}(\tau) = \gamma\big[\langle\hat{n}\rangle_0 - \mathcal{A}(\tau)\big]\Delta T + \mathcal{N}(0, \Sigma^2\Delta T), \tag{4.14}$$

where γ is the amplitude noise bandwidth and Σ^2 controls the strength of the random fluctuations. This noise model has a constant long-time ($\tau \to \infty$) variance of $\Sigma_\infty^2 = \Sigma^2/2\gamma$ as opposed to the phase noise, where the variance scales linearly with τ. The parameter $\langle\hat{n}\rangle_0$ is the mean value of the random walks in intensity and we also use this as the initial value such that $\langle\hat{n}\rangle_0 = \langle\mathcal{A}(\tau)\rangle = \mathcal{A}(0)$, where the average is taken across an ensemble of walks.

Implementing a Kalman filter for the phase and amplitude estimates improves the performance of the NN-based tracking method. This Kalman filter takes the raw estimates $(\hat{\mathcal{A}}_{NN}, \hat{\phi}_{NN})$ from the NN and outputs the final amplitude $\hat{\mathcal{A}}(\tau)$ and phase $\hat{\phi}$ estimates. Due to the finite number of experiments N per estimation, the NN estimator outputs have relatively large variance. The Kalman filter allows for incorporating this uncertainty along with the dynamics of the noise in an optimal way through the Kalman gain [203]. Algorithm 1 shows the full parameter tracking method based on the NN estimator with Kalman filtering. Initially, the LO power $\mathcal{B}(0)$ is set to the average input mean photon number $\langle\hat{n}\rangle_0$. The receiver then performs state discrimination measurements and adds the measurement data to the detection matrix $\mathbf{D}$ after each individual measurement. After completing $N = 10$ experiments, the receiver evaluates the NN to obtain the raw estimates $\hat{\phi}_{NN}$ and $\hat{\mathcal{A}}_{NN}$ for the phase offset and input power. These raw estimates are then passed to the Kalman filter where the filter incorporates the dynamics of the noise and returns updated,

Algorithm 1 NN Parameter Tracking Algorithm

Define: $x = 1 - \gamma\Delta T$
function KalmanPred(σ_ϕ^2, $\mathcal{B}(\tau)$, $\sigma_{\mathcal{A}}^2$)
  $\hat{y}_\phi \leftarrow 0$
  $\hat{\sigma}_\phi^2 \leftarrow \sigma_\phi^2 + N\sigma_1^2$
  $\hat{y}_{\mathcal{A}} \leftarrow x^N\mathcal{B}(\tau) + \langle\hat{n}\rangle_0(1 - x^N)$
  $\hat{\sigma}_{\mathcal{A}}^2 \leftarrow x^{2N}\sigma_{\mathcal{A}}^2 + \Sigma^2\Delta T(x^{2N} - 1)/(x^2 - 1)$
  return $\hat{y}_\phi, \hat{\sigma}_\phi^2, \hat{y}_{\mathcal{A}}, \hat{\sigma}_{\mathcal{A}}^2$
end function

function KalmanFilter($\mathcal{B}(\tau), \hat{A}_{NN}, \hat{\phi}_{NN}$)
  $\hat{y}_\phi, \hat{\sigma}_\phi^2, \hat{y}_{\mathcal{A}}, \hat{\sigma}_{\mathcal{A}}^2 \leftarrow$ KalmanPred(σ_ϕ^2, $\mathcal{B}_N$, $\sigma_{\mathcal{A}}^2$)
  $K_\phi \leftarrow \hat{\sigma}_\phi^2/(\hat{\sigma}_\phi^2 + \sigma_{\phi,NN}^2)$ ▷ Phase Kalman gain
  $\hat{\phi} \leftarrow K_\phi\hat{\phi}_{NN} + (1 - K_\phi)\hat{y}_\phi$ ▷ Phase estimate
  $\sigma_\phi^2 \leftarrow (1 - K_\phi)\hat{\sigma}_\phi^2$ ▷ Phase variance
  $K_{\mathcal{A}} \leftarrow \hat{\sigma}_{\mathcal{A}}^2/(\hat{\sigma}_{\mathcal{A}}^2 + \sigma_{\mathcal{A},NN}^2)$ ▷ Amp. Kalman gain
  $\hat{\mathcal{A}} \leftarrow K_{\mathcal{A}}\hat{A}_{NN} + (1 - K_{\mathcal{A}})\hat{y}_{\mathcal{A}}$ ▷ Amp. estimate
  $\sigma_{\mathcal{A}}^2 \leftarrow (1 - K_{\mathcal{A}})\hat{\sigma}_{\mathcal{A}}^2$ ▷ Amp. variance
  return $\hat{\phi}, \hat{\mathcal{A}}(\tau)$
end function

Initial: $\mathcal{B}(0) \leftarrow \langle\hat{n}\rangle_0$, $\delta(0) \leftarrow 0$
$\tau \leftarrow 0$ ▷ Time in increments of symbol time
$n \leftarrow 0$ ▷ Number of measurements performed
loop
  Do Single state discrimination measurement
  $\{d_j\}_L, \{\Delta_j\}_L \leftarrow$ Single measurement data
  Add $\{d_j\}_L, \{\Delta_j\}_L$ to detection matrix $\mathbf{D}$
  $\tau \leftarrow \tau + \Delta T$
  $n \leftarrow n + 1$
  if $n = N$ **then** ▷ Update LO
    $\hat{A}_{NN}, \hat{\phi}_{NN} \leftarrow$ Evaluate NN
    Reset $\mathbf{D}$ to zeros
    $\hat{\phi}, \hat{\mathcal{A}}(\tau) \leftarrow$ KalmanFilter($\mathcal{B}(\tau), \hat{A}_{NN}, \hat{\phi}_{NN}$)
    $\mathcal{B}(\tau) \leftarrow \hat{\mathcal{A}}(\tau)$ ▷ Correct LO intensity
    $\delta(\tau) \leftarrow \delta(\tau) + \hat{\phi}$ ▷ Correct LO phase
    $n \leftarrow 0$ ▷ Reset measurement counter
  else ▷ Don't update LO
    $\mathcal{B}(\tau) \leftarrow \mathcal{B}(\tau - \Delta T)$
    $\delta(\tau) \leftarrow \delta(\tau - \Delta T)$
  end if
end loop

filtered estimates $\hat{\phi}, \hat{\mathcal{A}}(\tau)$ of the phase and amplitude noise. Finally, these filtered estimates are used to correct the state discrimination measurement by adjusting the LO power to $\mathcal{B}(\tau) = \hat{\mathcal{A}}(\tau)$, and the LO phase by $\delta(\tau) = \delta(\tau - \Delta T) + \hat{\phi}$ such that $\arg(\beta) = \theta_j + \delta(\tau)$. We obtain the prediction equations in the KalmanPred$(\cdot)$ function by first propagating the stochastic differential equations for the phase (Eq. (4.13)) and amplitude (Eq. (4.14)) noise forward in time by N time steps and then taking the mean and variance.

4.2.3 Simulation Results

We benchmark the performance of the NN-based tracking method against an approach based on a Bayesian estimator, which has a more complex implementation than the NN estimator. Figure 4.12(a) shows the simulated probability of error for an adaptive non-Gaussian receiver implementing the NN-based noise tracking method as a function

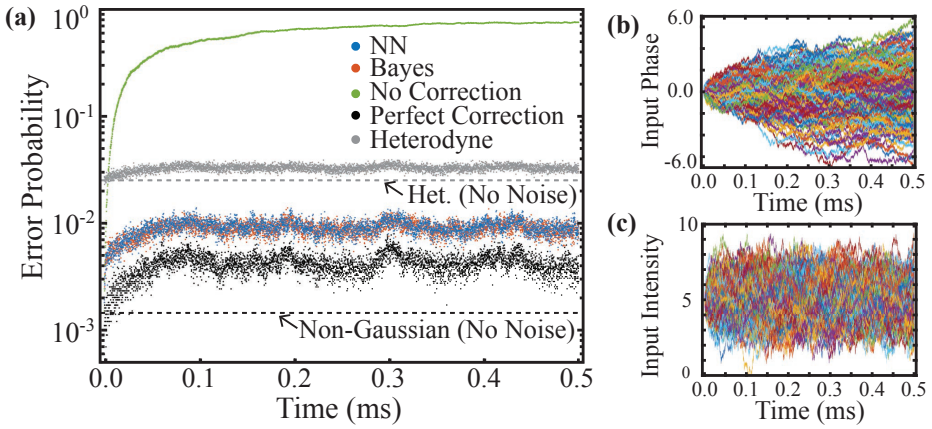

Figure 4.12: (a) Error probability for the adaptive non-Gaussian receiver in the presence of channel noise. The green points show an uncorrected measurement while the blue and orange points show the results for a NN-based and Bayesian-based tracking method. The black points show the error for a perfectly corrected non-Gaussian measurement and the gray points show the error for a heterodyne receiver with perfect correction. (b,c) Waveforms of the applied (b) phase noise at $\Delta\nu = 5$ kHz and (c) amplitude noise with $\gamma = 25$ kHz and $\Sigma^2_\infty = 1.5$. Figure adapted from: *Physical Review Research* 3, 013200 (2021).

of time with $\langle \hat{n} \rangle_0 = 5.0$ and $N = 10$. We simulate 1000 different realizations of phase noise (b) with $\Delta\nu = 5$ kHz, and amplitude noise (c) with $\gamma = 25$ kHz and $\Sigma^2_\infty = 1.5$. The blue and orange points show the results of the noise tracking using methods based on a NN and Bayesian estimator, respectively. The black points show the results if the receiver implements perfect noise tracking, which refers to the receiver having complete knowledge of the noise such that $\mathcal{B}(\tau) = \mathcal{A}(\tau)$ and $\delta(\tau) = \phi(\tau)$. The green points show the error probability when the receiver does not implement any noise tracking and the gray points show the performance of an ideal heterodyne with perfect phase tracking, which serves as the equivalent QNL in the presence of noise. The dashed black and gray lines show the error probability for the heterodyne and non-Gaussian receivers in the absence of noise. We find that the NN-based method achieves the same performance as a more complex Bayesian estimator, and both methods can maintain an error probability below the limit of ideal heterodyne detection.

We study the performance of the NN-based method for parameter tracking with different phase and amplitude noise parameters to demonstrate the versatility of this

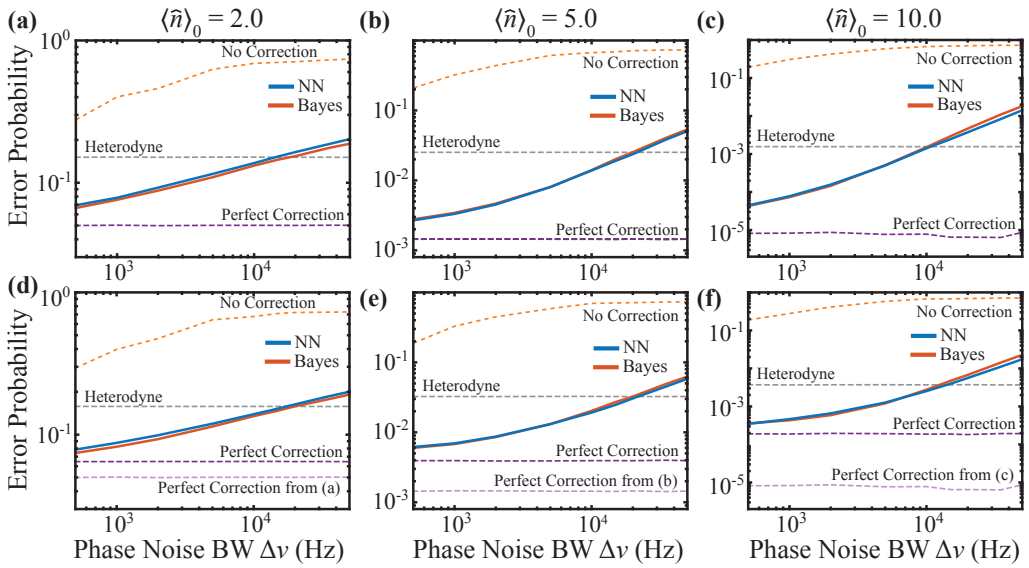

Figure 4.13: Results for the error probability as a function of the phase noise bandwidth $\Delta\nu$ without (a-c) and with (d-f) amplitude noise at $\gamma = 25$ kHz, Σ^2_∞ =0.25, 1.5, and 6.0 for an average input power of $\langle\hat{n}\rangle_0 = 2.0$, 5.0, and 10.0, respectively. The blue and orange lines show the error for the NN and Bayesian based tracking methods. The gray dashed lines show a perfectly corrected heterodyne receiver and the purple dashed lines show a perfectly corrected non-Gaussian measurement. Figure adapted from: *Physical Review Research* 3, 013200 (2021).

method. We note that for all simulations, we fix the number of measurements per estimation to $N = 10$. Figure 4.13 shows the performance of the NN-based tracking method (solid blue) compared to noise tracking method based on a Bayesian estimator (solid orange) as a function of phase noise bandwidth $\Delta\nu$ for an average input energy of $\langle\hat{n}\rangle_0 = 2.0$, 5.0, and 10.0. Panels (a)-(c) show the results in the absence of amplitude noise ($\Sigma^2_\infty = 0$, $\gamma = 0$), corresponding to the one dimensional problem of phase tracking described in Ch. 4.1. Panels (d)-(f) show the performance in the presence of both phase noise and amplitude noise at γ=25 kHz with Σ^2_∞=0.25, 1.5, and 6.0, respectively. In all panels, the dashed lines show the results for an ideal heterodyne receiver with perfect tracking (gray) and a non-Gaussian receiver implementing perfect tracking (purple). We find that even in the presence of severe amplitude noise (lower panels), the NN-based parameter tracking method maintains performance equivalent to the Bayesian approach. The NN-based method enables sub-QNL performance up

to a phase noise bandwidth of $\Delta\nu \approx$10 kHz for all input powers studied. In some situations (b,c,e,f), the NN-based method appears to slightly outperform the Bayesian approach. This result is potentially due to the fact that a Bayesian estimator may not be the most precise estimator given the limited number of samples N per estimate, as opposed to the asymptotic optimal behavior [204]. Furthermore, the relative weight of the error in phase estimates and error in amplitude estimates can be customized in the training process of the NN estimator. This freedom may allow for the NN-based method to better minimize the overall error probability.

We also investigate the effect of different amplitude noise bandwidths γ with fixed strength Σ^2_∞ in the presence of constant phase noise with bandwidth $\Delta\nu$. Figure 4.14(a-c) shows the probability of error as a function of the amplitude noise γ for a fixed phase noise level $\Delta\nu$=5 kHz for $\langle\hat{n}\rangle_0$=2.0, 5.0, and 10.0, respectively. The blue and orange lines shows the results for the NN-based and Bayesian-based tracking method, respectively, both with ($\Delta\nu$=5 kHz) and without ($\Delta\nu = 0$) phase noise. We find that the NN-based method again performs similar to the Bayesian method across a wide range of γ and that increasing γ has a less severe effect of the overall error compared to increasing the phase noise bandwidth.

Figure 4.14(a) shows that both the NN and Bayesian based methods achieve error rates below that of a perfectly corrected measurement for $\Delta\nu = 0$. At first glance, this should not be possible since perfect correction corresponds to the receiver

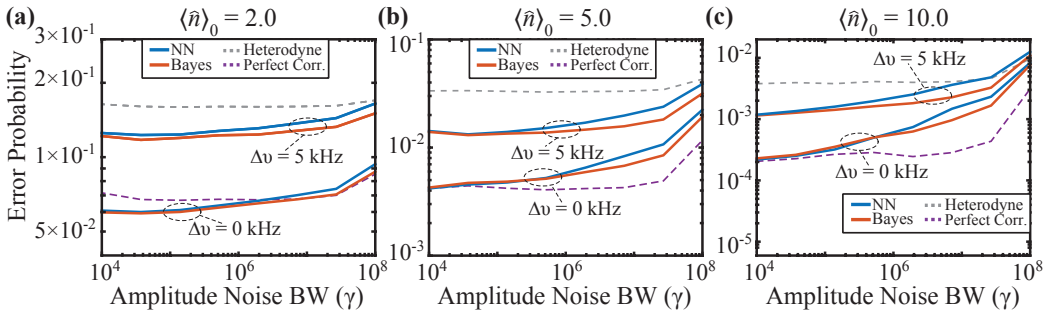

Figure 4.14: (a-c) Error probability as a function of amplitude noise bandwidth γ both without and with phase noise at $\Delta\nu$ =5 kHz with Σ^2_∞=0.5, 1.5, and 6.0 for $\langle\hat{n}\rangle_0$ = 2.0, 5.0, and 10.0, respectively. Figure adapted from: *Physical Review Research* 3, 013200 (2021).

having complete knowledge of the time dependent noise. However, we find that this behavior is due to an underlying bias in the amplitude estimates when $\langle \hat{n} \rangle_0 = 2.0$. Both the NN and Bayesian estimators are biased towards estimates which are larger than the actual input energy, causing the corrected LO power to be biased as well such that $\mathcal{B}(\tau) > \mathcal{A}(\tau)$. This bias results in the corrected receiver incidentally approximating an optimized measurement strategy for state discrimination similar to the "flat optimization" strategy described in Sec. 3.2. As discussed in Sec. 3.2, optimized measurements for state discrimination can achieve error rates which are below that of a nulling strategy ($\mathcal{B}(\tau) = \mathcal{A}(\tau)$), which corresponds to perfect correction here. Thus, the overall effect of the estimator bias is that both parameter tracking methods can actually perform slightly better than a perfectly corrected receiver, albeit by accident. Note that understanding how to implement parameter tracking with optimized measurements is an interesting topic of further study due to their applications in the single-photon regime.

4.2.4 Strategy Comparison

While previously we fixed $N = 10$ for the simulations to demonstrate the versatility and robustness of this approach, an optimal value for N can be found for specific channel noise levels as in Sec. 4.1.5. If the channel noise parameters are well known, then the optimal value of N minimizes the overall error probability, which depends on the total variance σ_{tot}^2 of the estimates. This variance is comprised of the intrinsic variance of the estimator $(\Delta\hat{\phi})^2$ and the noise variance, both of which are dependent on the value of N. Increasing the value of N will reduce the estimator variance due to the statistical $1/N$ scaling while the noise variance will increase linearly with N. On the other hand, decreasing N causes the intrinsic estimator variance to dominate the total variance. While the Kalman filter mainly takes care of finding the optimal balance between these two uncertainties, there is still a need for selecting an optimal value of N which achieves the best performance.

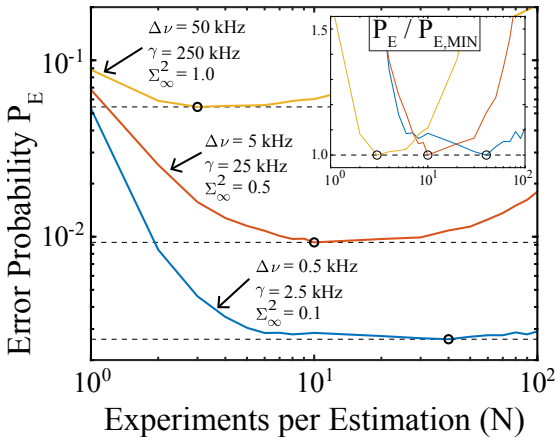

Figure 4.15: Error probability at $\langle \hat{n} \rangle = 5.0$ for small (blue), moderate (orange), and severe (yellow) channel noise as a function of N. The black points each show the optimal values of N and the respective error probability. The inset shows the error for each curve normalized by the minimum value to easily compare the optimal N for each noise level. Figure adapted from: *Physical Review Research* 3, 013200 (2021).

Figure 4.15 shows the error probability for a $\langle \hat{n} \rangle_0 = 5.0$ strategy as a function of N for three different noise regimes. The blue, orange, and yellow lines show the error for small ($\Delta\nu = 0.5$ kHz, $\gamma = 2.5$ kHz, $\Sigma^2_\infty = 0.1$), moderate ($\Delta\nu = 5$ kHz, $\gamma = 25$ kHz, $\Sigma^2_\infty = 0.5$), and severe ($\Delta\nu = 50$ kHz, $\gamma = 250$ kHz, $\Sigma^2_\infty = 1.$) channel noise, respectively. The optimal number of experiments per estimation for low levels of noise is $N \approx 40$, while for high noise levels the optimal value is $N \approx 3$. The inset shows the error probability normalized by the minimum error for each noise level. We find that $N = 10$ is optimal for moderate noise levels, and the error at $N = 10$ for the other noise regimes is only slightly higher than the respective minima. This finding shows that while there is optimal value of N for a specific level of channel noise, a different value of N may be ideal for robust and versatile implementations.

4.2.5 Discussion

Neural networks can implement highly accurate parameter estimation while being far more computationally efficient compared to a Bayesian estimator. Here, the neural network performs tracking of phase and amplitude channel noise but can

easily scale to higher dimensions. The neural network based algorithm maintains sub-QNL error rates across a wide range of input powers, noise strengths and noise bandwidths. The network can even be tailored to devices such as FPGAs by tailoring the architecture to allow for a high degree of parallelization, and using non-linear activation functions with simple implementations. Furthermore, this NN-based method can potentially be applied to optimized state discrimination strategies in the single-photon regime to facilitate quantum communication protocols, which rely on accurate state discrimination with low input powers.

4.3 Conclusion

In this chapter, we experimentally demonstrated a phase tracking algorithm for adaptive non-Gaussian receivers based on a simple and efficiently calculable estimator. We implemented this phase tracking method for different input powers and noise strengths and investigated the trade-off between estimation accuracy and bandwidth as a function of the resources used for estimation [85]. As an extension of phase tracking to multi-dimensional channel noise, we also developed an algorithm for simultaneous tracking of phase and amplitude noise based on a neural network. Through simulations we show that this method based on a neural network can potentially enable non-Gaussian receivers to operate below the QNL across a wide range of channel noise parameters. This algorithm can in principle be implemented at very high bandwidth due to its computational efficiency compared to a Bayesian estimator, and can easily scale to even higher dimensional estimation [86]. Efficient and accurate channel noise tracking is a critical barrier to practical implementations of adaptive photon counting measurements and the techniques explored in this chapter will potentially enable transitioning these technologies out of the lab.

Chapter 5

Phase Estimation of Coherent States

5.1 Introduction

Estimation of a single parameter, in particular that of optical phase estimation, is a central problem in physics and metrology [12, 205, 206]. In its simplest form, a field of light (probe state) first acquires a phase shift ϕ through interaction with a system. A measurement of the light then results in an estimate $\hat{\phi}$ for the unknown phase ϕ. The goal of any measurement of course, is to achieve the smallest possible uncertainty $(\Delta\hat{\phi})^2$ in estimation, and any measurement will fundamentally only be able to extract as much information as the probe state can possibly carry [1, 4, 7, 10, 207–210]. The quantum Fisher information (QFI) quantifies how much information a given probe state can contain [6, 7, 211, 212] and is a useful tool for understanding the capabilities of different states of light. The QFI also bounds the attainable uncertainty of estimation through the Cramer-Rao lower bound (CRLB) [6]: $(\Delta\phi)^2 \geq 1/(M \times \text{QFI})$, where M is the number of repetitions of the experiment. The CRLB is proportional to $1/N$ when using coherent states of light, where N is the amount of energy contained in the probe state [5]. This limit on the uncertainty when using coherent states is also referred to as

the standard quantum limit (SQL). On the other hand, states of light with quantum correlations can allow for the CRLB to scale as $1/N^2$, i.e. Heisenberg scaling, and quantum states of light can be used to achieve sub-SQL uncertainties [47, 213–219]. Phase estimation beyond the SQL using light with quantum correlations has been investigated [177, 178, 220–223] for sensing small phase deviations from a known phase [213, 224–229], with repeated sampling ($M > 1$) [230, 231], and with feedback measurements [177, 178, 220, 221, 224].

Here, we focus on using coherent states of light as opposed to states with quantum correlations or entangled states. Phase estimation of coherent states is essential for mechanical cooling of oscillators [232–234], preparation of spin squeezed atomic states [235], as well as force and waveform detection [47, 236]. Furthermore, we study "single-shot" phase estimation where $M = 1$, a challenging situation which allows for exploring how to extract the most amount of information possible in a single measurement. Much attention has been paid to adaptive homodyne measurements for single-shot phase estimation of coherent states, which have been shown in theory to asymptotically approach the CRLB [78, 237–241]. Experimentally, adaptive homodyne measurements [242] can outperform an equivalent (adjusted for efficiency) heterodyne receiver, whose uncertainty is bounded by *twice* the CRLB for coherent states. On the other hand, photon counting based measurements have shown that non-Gaussian approaches have the capability to surpass the limits of heterodyne detection [228, 243]. However, the limits of adaptive non-Gaussian measurements for single-shot phase estimation of coherent states is an open question.

In this chapter, we develop and implement adaptive non-Gaussian strategies for estimation of an unknown phase of a coherent state. These strategies are based on optimized displacement operations, photon counting and feedback. We find that real-time optimization of the displacement operation conditioned on the measurement record allows for surpassing the limit of ideal heterodyne detection without correcting for any inefficiencies [87]. We develop a Bayesian updating algorithm and optimization strategy based on a Gaussian approximation which can be experimentally implemented

at relatively high bandwidth. We study two different optimization strategies and find that while on the surface the two strategies seem very different, they achieve almost identical results and can potentially outperform the best known adaptive homodyne measurement [88].

5.2 Phase Estimation Strategy

The strategy for phase estimation of coherent states of light relies on displacement operations, photon counting, feedback, and real-time optimization of the displacements conditioned on the measurement record [87]. Figure 5.1 shows a schematic of the phase estimation strategy, which follows a recursive Bayesian approach consisting of L adaptive steps. At the beginning of a single adaptive step j, the strategy first obtains a value β_j for the displacement operation $\hat{D}(\beta_j)$. A local oscillator (LO) then displaces the input state $|\alpha e^{i\phi_0}\rangle$ with unknown phase ϕ_0 in phase space to the state $\hat{D}(\beta_j)|\alpha e^{i\phi_0}\rangle = |\alpha e^{i\phi} - \beta_j\rangle$, which is detected by a single photon detector (SPD) with

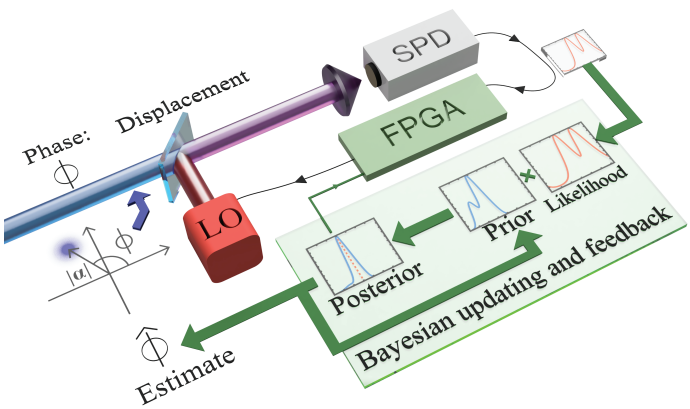

Figure 5.1: Diagram of the adaptive strategy for phase estimation of coherent states with real-time optimization. In each adaptive step, the input state $|\alpha e^{i\phi}\rangle$ is displaced $\hat{D}(\beta_{opt})$ by the optimized LO field and detected by a single photon detector (SPD) with photon number resolution. The strategy calculates the posterior distribution for the adaptive step conditioned on the detection result and β_{opt}. The LO is then re-optimized for the next adaptive step given the updated phase distribution. Figure adapted from: *Physical Review Letters* 125, 120505 (2020).

PNR(m). The strategy updates the posterior distribution $P(\phi|d_j, \beta_j)$ for the phase using Bayes rule given a particular measurement outcome d_j and displacement β_j:

$$P(\phi|d_j, \beta_j) = \frac{1}{P(d_j)}\mathcal{L}(d_j|\phi, \beta_j)P(\phi). \tag{5.1}$$

Here $\mathcal{L}(d_j|\phi, \beta_j)$ is the likelihood function, $P(\phi)$ is the prior probability distribution, and $P(d_j) = \int \mathcal{L}(d_j|\phi, \beta_j)P(\phi)d\phi$ is the total probability of detecting d_j photons. The posterior distribution $P(\phi|d_j, \beta_j)$ from step j then becomes the prior probability distribution $P(\phi)$ for the next adaptive step $j + 1$.

The strategy obtains the displacement value β_j at the beginning each adaptive step based on the current prior probability distribution $P(\phi)$. If the displacement $\hat{D}(\beta_j)$ is set to displace the most likely state to the vacuum state, as in state discrimination measurement strategies, then the estimation variance cannot reach below the limit of ideal heterodyne detection [87, 88]. Instead, the receiver optimizes the displacement β_j at the beginning of every adaptive step given $P(\phi)$, which yields an optimal value β_{opt} given the entire detection history $\{d_1, ..., d_{j-1}\}$ up to the current adaptive step j. Optimization of the displacement operation requires a particular objective function, and we implement strategies based on two different functions [87, 88]. One possible objective function is the average sharpness of the posterior distribution, which has been explored in the context of adaptive Gaussian measurements [223, 244] and interferometry [177, 224]. The average sharpness $\langle S(\beta, m)\rangle$ is given by:

$$\langle S(\beta, m)\rangle = \sum_{n=0}^{m} P(n)\left|\int_0^{2\pi} e^{i\phi}P(\phi|n, \beta)\right| = \sum_{n=0}^{m}\left|\int_0^{2\pi} e^{i\phi}\mathcal{L}(n|\phi, \beta)P(\phi)\right|, \tag{5.2}$$

where the average is taken over all possible detection results for adaptive step j and the second equality is obtained by using Bayes theorem. If the sharpness of a distribution is large, then the spread is small and vice versa, and optimization yields a displacement value β_{opt} which maximizes $\langle S(\beta, m)\rangle$.

The second objective function we study is the mutual information [241, 245, 246]:

$$I(\beta, m) = \sum_{n=0}^{m}\int_0^{2\pi} P(\phi, n|\beta)\log_2\left(\frac{P(\phi, n|\beta)}{P(n)P(\phi)}\right)d\phi, \tag{5.3}$$

where $P(\phi, n|\beta) = P(\phi|n,\beta)P(n) = \mathcal{L}(n|\phi,\beta)P(\phi)$ is the joint probability distribution for ϕ and n. The mutual information can also be written in terms of the Kullback-Liebler divergence, or relative entropy [247, 248]:

$$I(\beta, m) = \sum_{n=0}^{m} P(n) D_{KL}(P(\phi|n,\beta)||P(\phi)) \tag{5.4}$$

$$= \int_0^{2\pi} P(\phi) D_{KL}(\mathcal{L}(n|\phi,\beta)||P(n))d\phi, \tag{5.5}$$

where the Kullback-Liebler divergence $D_{KL}(P(x)||Q(x))$ between two distributions $P(x)$ and $Q(x)$ is given by:

$$D_{KL}(P(x)||Q(x)) = \int P(x)\log_2\left(\frac{P(x)}{Q(x)}\right)dx. \tag{5.6}$$

Therefore, maximization of the mutual information will yield an optimal displacement β_{opt} which maximizes the average relative entropy between the posterior and prior distributions, averaged over possible photon detections.

We optimize the same objective function each adaptive step such that there are two different measurement strategies, one which maximizes the sharpness and one which maximizes the mutual information. Application of this recursive Bayesian updating and optimization over all L adaptive steps for a single input pulse results in a history of detections $\{d_j\}_L$ and optimal displacement values $\{\beta_{opt}\}_L$. We calculate the final phase estimate given the measurement data as [6]:

$$\hat{\phi} = \arg\left(\int_0^{2\pi} e^{i\phi} P(\phi|\{d_j\}_L, \{\beta_{opt}\}_L)d\phi\right), \tag{5.7}$$

where $P(\phi|\{d_j\}_L, \{\beta_{opt}\}_L)$ is the reconstructed posterior distribution given the entire measurement record. Figure 5.2 shows a flowchart of the adaptive phase estimation strategy with optimized displacement operations, which starts with an initially uniform prior probability distribution for the input phase. The strategy uses the prior probability distribution for each adaptive step to optimize the displacement operation and then calculates the posterior probability distribution given a particular

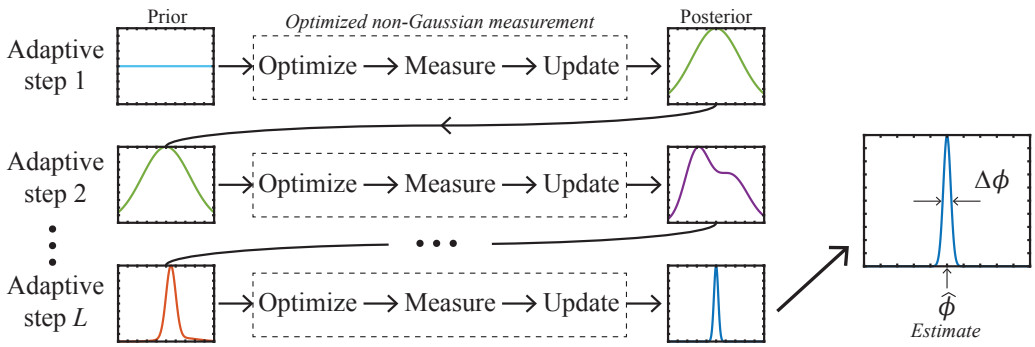

Figure 5.2: Flowchart of the adaptive measurement strategy based on optimized displacement operations and feedback. The belief for the input phase is maintained as a probability distribution and updated according to Bayes rule. The posterior probability distribution for a given adaptive step becomes the prior probability distribution for the next adaptive step.

measurement result. The posterior distribution then becomes the prior probability distribution for the next adaptive step and the strategy uses this new, updated prior distribution to re-optimize the displacement in a recursive Bayesian approach.

This adaptive measurement scheme is repeated many (N) times, each with a random initial relative phase between the input and LO. We then calculate the variance $Var[\hat{\phi}]$ of the distribution of the N final phase estimates via the Holevo variance for cyclic variables [6, 240, 249], which is bounded by the CRLB:

$$Var[\hat{\phi}] = \frac{1}{|\langle e^{i\hat{\phi}}\rangle|^2} - 1 \geq \frac{1}{4|\alpha|^2}, \tag{5.8}$$

where $|\langle e^{i\hat{\phi}}\rangle| = |\sum_{k=1}^{N} e^{i\hat{\phi}_k}|/N$ corresponds to the sharpness of the distribution of final phase estimates and $4|\alpha|^2$ is the QFI for coherent states [2, 6, 7]. This Holevo variance characterizes the estimation uncertainty for cyclic variables, such as the phase, as opposed to the typical mean-squared error for variables which are not cyclic. However, the mean-squared error and Holevo variance coincide in the limit that the distribution of final phase estimates has small uncertainty [6, 9].

Figure 5.3 shows the theoretical expectations of the optimized non-Gaussian strategy for phase estimation with L=30 adaptive steps, PNR(3), ideal detection efficiency $\eta = 1.0$, and imperfect visibility $\xi = 0.997$. The orange and blue solid lines show the results for strategies which optimize the mutual information and sharpness,

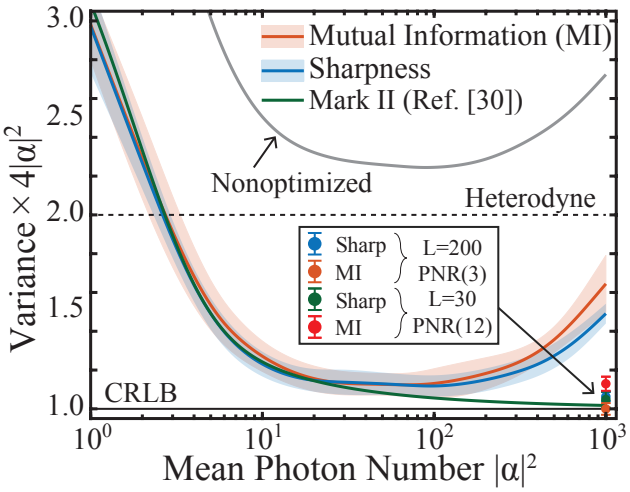

Figure 5.3: Expected performance of the optimized strategies for phase estimation of coherent states. The orange and blue lines show the results for strategies which optimize the mutual information and sharpness, respectively, every adaptive step. The solid lines and shaded regions show the average and standard deviation for five sets of 10^3 independent Monte Carlo samples. The gray line shows a non-optimized strategy and the green line shows the variance of the "Mark II" adaptive homodyne strategy from Ref. [78]. The black solid and dashed lines show the CRLB for coherent states and limit of an ideal heterodyne measurement, respectively. The points and error bars at $|\alpha|^2 = 10^3$ show both strategies for $L = 200$ with PNR(3) and $L = 30$ with PNR(12) for comparison. Figure adapted from: *Physical Review Letters* 125, 120505 (2020).

respectively. The shaded regions correspond to one standard deviation across five data sets of $N = 10^3$ individual experiments each. The black solid line shows the CRLB for coherent states $(1/4|\alpha|^2)$ and the black dashed line shows the lower bound on an ideal heterodyne measurement $(1/\mathrm{CFI}_{het})$, where $\mathrm{CFI}_{het} = 2|\alpha|^2$ is the classical Fisher information for a heterodyne receiver. The green solid line shows the expected results for the best known adaptive homodyne strategy termed "Mark II" proposed in Refs. [78, 240] and the gray line shows the result for a non-optimized strategy. We find that both optimization strategies perform almost equivalently except for very large mean photon numbers and reach a minimum variance of ≈ 1.1 times the CRLB at $|\alpha|^2 \approx 100$. Furthermore, both strategies can achieve the same performance as the "Mark II" measurement for $|\alpha|^2 \lessapprox 25$. The increase in the variance at high input powers is mainly a result of the limited number of adaptive steps and PNR. Increasing either the PNR or the number of adaptive steps allows the measurement

to reach performance very close to the CRLB [87, 88]. The blue and orange points show the variance of the sharpness and mutual information strategies for L=200 with PNR(3), respectively, and the green and red points show both strategies for L=30 with PNR(12). The strategy maximizing the mutual information with L=200 and PNR(3) achieves the smallest variance of 1.003 ± 0.03 times the CRLB. Intuitively, if the mean photon number per adaptive step ($|\alpha|^2/L$) is much larger than the photon number resolution m, then the strategy will not be able to achieve small variances. If the ratio $|\alpha|^2/(mL)$ increases, the measurements ability to fully resolve the state decreases, leading to an increase in the measured variance. However, this effect at larger values of $|\alpha|^2$ can be avoided by increasing either L or the PNR [88].

5.3 Experimental Implementation & Real-Time Optimization Algorithm

We use the interferometric setup described in Appendix A.2 to experimentally implement single-shot phase estimation of optical coherent states. Figure 5.4 shows a schematic of the experimental setup. The main difference from previous experiments is an upgrade in the field-programmable gate array (FPGA) to an Altera Cyclone IV (EP4CE55F23C8N), which has approximately an order of magnitude more logical elements and memory blocks compared to previously used FPGAs. We also now can continuously control both the LO phase and amplitude using an 8-bit digital-to-analog converter (DAC) for each connected to the FPGA. A Helium-Neon laser and pulsed acousto-optic modulator (AOM) generate $\approx 600\ \mu s$ optical pulses and an attenuator (Att.) and phase modulator (PM) prepare the input coherent state with mean photon number $|\alpha|^2$. For experimental convenience, instead of randomly choosing the initial input phase ϕ_0 with a fixed initial phase of the LO $\theta_0 = 0$, we equivalently fix $\phi_0 = \pi$ and start the LO at a random phase θ_0 such that the initial *relative* phase $\phi_0 - \theta_0$ is still random. An amplitude modulator (AM) and second PM prepare the LO field used to implement the displacement $\hat{D}(\beta_{opt})$ via interference on a 99/1 beam

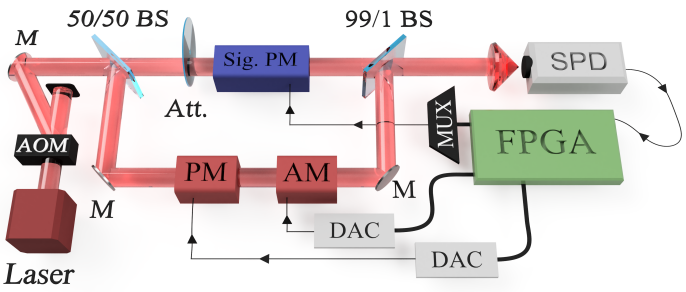

Figure 5.4: Experimental setup for phase estimation of coherent states. Compared to previous experimental setups, both the phase and amplitude of the LO field used for the displacement operation are fully controllable using digital-to-analog converters (DACs). An upgraded field-programmable gate array (FPGA) implements the phase estimation algorithm including the the Bayesian updating as well as the approximate optimization. Figure adapted from: *Physical Review Letters* 125, 120505 (2020).

splitter [79]. A single photon detector (SPD) then detects the displaced state with photon number resolution (See Appendix B.1 for details). The FPGA implements the Bayesian updating procedure ($\approx$ 500 ns) as well as the real-time optimization ($\approx$ 500 ns) of the displacement every adaptive step. Our implementation uses $L = 30$ adaptive steps of ≈ 20 μs each with PNR(3) and achieves an overall system efficiency of $\eta = 0.70$, interference visibility of $\xi = 0.997$, and dark counts of $\nu = 10^{-3}$ per experiment.

The phase estimation measurement relies on obtaining an optimal displacement value β_{opt} conditioned on the detection history. We denote the full prior probability distribution for adaptive step j as $P^{(j)}(\phi)$ and the posterior distribution given detection d_j and β_{opt} as $P^{(j)}_{post}(\phi) = P(\phi|d_j, \beta_{opt})$ from Eq. (5.1). Optimization of the sharpness $\langle S(\beta, m)\rangle$ in Eq. (5.2) or mutual information $I(\beta, m)$ in Eq. (5.5) given $P^{(j)}(\phi)$ is simple to perform in MATLAB. However, the experimental implementation requires very high-speed calculation of β_{opt} on the order of hundreds of nanoseconds. To this end, we develop a real-time *approximate* optimization strategy for quickly obtaining β_{opt} at each adaptive step j in order to experimentally implement the strategy in the FPGA. We decompose the phase of the LO $\theta_j = \arg(\beta_j)$ into the current maximum *a posteriori*

probability (MAP) estimate $\hat{\phi}_{MAP} = \text{argmax}\Big(P_{post}^{(j-1)}(\phi)\Big)$ plus an optimized offset Δ_{opt} such that $\arg(\beta_{opt}) = \theta_{opt} = \hat{\phi}_{MAP} + \Delta_{opt}$. The phase range $\phi = [0, 2\pi]$ is discretized into 256 phase points and we use 10 bits of precision to represent each probability $P^{(j)}(\phi)$ in the FPGA with fixed point notation, i.e. the bit-string $\{b_i[0], b_i[1], ..., b_i[9]\}$ denotes a single probability as $P^{(j)}(\phi_i) = \sum_{k=0}^{9} 2^{-b_i[k]}$ with $b[k] \in \{0, 1\}$.

Figure 5.5 shows a flowchart of the algorithm for phase estimation implemented by the FPGA. Parts (I-III) show how the algorithm splits finding the optimal magnitude $|\beta|_{opt}$ and offset Δ_{opt} into two cases depending on the adaptive step j:

- **Case I ($j \leq 8$)** - The strategy obtains the optimal magnitude $|\beta|_{opt}$ and offset Δ_{opt} from a pre-calculated look-up-table (LUT) whose input is the entire detection record $\{d_1, d_2, ..., d_{j-1}\}$ up to adaptive step j. We construct this LUT offline by simulating all possible detection histories and maximizing either the sharpness or mutual information using the full prior probability distribution $P^{(j)}(\phi)$ to obtain the exactly optimal values for $|\beta|_{opt}$ and Δ_{opt}.

- **Case I ($j \geq 9$)** - The strategy approximates the full prior distribution $P^{(j)}(\phi)$ for adaptive step j as a skewed Gaussian distribution for the purposes of efficient real-time calculation of *approximately* optimal values for $|\beta|_{opt}$ and Δ_{opt}. The FPGA calculates the variance σ^2 and skewness μ_3 of $P^{(j)}(\phi)$ (Step II) and inputs the variance σ^2 to a second LUT, which outputs an approximately optimal LO magnitude $|\beta|_{opt}$ and value for the phase offset δ (Step III). We pre-calculate this second LUT offline by approximating the prior distribution as a zero-mean Gaussian distribution with variance σ^2 and calculating $|\beta|_{opt}$ and δ under that approximation (see Sec. 5.5). The algorithm obtains the approximately optimal phase offset $\Delta_{opt} = f(\mu_3) \times \delta$ using the customized function $f(\mu_3) = \text{sign}(\mu_3)/(1 + |\mu_3|)$.

We obtain the function $f(\mu_3)$ empirically based on Monte Carlo simulations by comparing the optimal phase offsets obtained using the Gaussian approximation and the full prior probability distribution. This Gaussian approximation is justified by

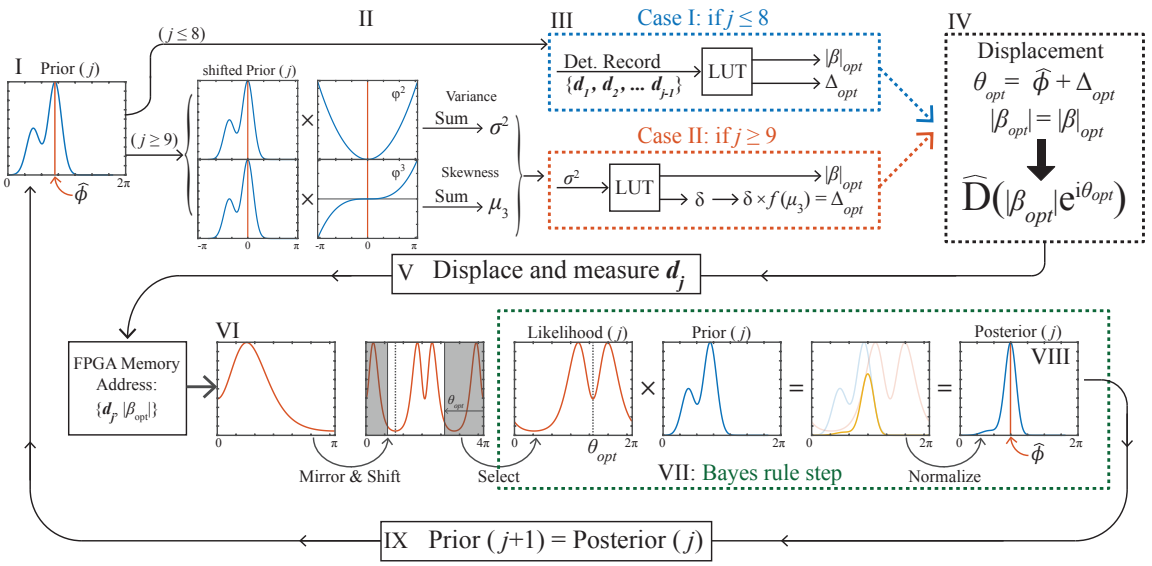

Figure 5.5: Flowchart of the phase estimation algorithm based on real-time optimization of the displacement operation. At the beginning of each adaptive step j, the algorithm obtains optimal LO values $\Delta_{opt}, |\beta|_{opt}$ either exactly if $j \leq 8$ or approximately if $j \geq 9$. The optimized displacement $\hat{D}(|\beta|_{opt}e^{i\theta_{opt}})$ is then implemented and the detected number of photons d_j is used to calculate the posterior distribution given the measurement outcome. Figure adapted from: *Physical Review Letters* 125, 120505 (2020).

observing that beyond $j \approx 9$, $P^{(j)}(\phi)$ is well described by its first three moments: the mean ($\approx \hat{\phi}_{MAP}$), variance σ^2, and skewness μ_3, and that the support over higher order moments is negligible. Steps IV and V are the implementation of the displacement $\hat{D}(\beta_{opt})$ with $\beta_{opt} = |\beta|_{opt}e^{i(\hat{\phi}_{MAP}+\Delta_{opt})}$ followed by detection of d_j photons.

Steps VI through IX implement the Bayesian updating of the distribution for the unknown phase. Step VI reconstructs the complete likelihood function $\mathcal{L}(d_j|\phi, \beta_{opt})$ given a particular detection d_j, displacement magnitude $|\beta|_{opt}$, and displacement phase θ_{opt}. Storing the likelihood function $\mathcal{L}(d_j|\phi, \beta_{opt})$ for every possible displacement phase and magnitude for each detection in the FPGA memory would require too many resources. Instead, we store a basic piece of the likelihood functions for photon counting from which the algorithm reconstructs the full likelihood function $\mathcal{L}(d_j|\phi, \beta_{opt})$. This reconstruction procedure is described in detail in Appendix B.4 and is based on the symmetries and circular properties of the likelihood functions for photon counting.

Step VII (green box) implements the Bayesian updating for adaptive step j where the FPGA multiplies the reconstructed likelihood function $\mathcal{L}(d_j|\phi, \beta_{opt})$ by the prior probability distribution $P^{(j)}(\phi)$ to obtain the updated but un-normalized posterior distribution $P^{(j)}_{post}(\phi)$. Step VIII updates the phase estimate $\hat{\phi}_{MAP}$ as the phase with maximum posterior probability, which is done using $P^{(j)}_{post}(\phi)$ since the MAP estimate is not a function of the normalization factor. Step VIII also implements a quasi-normalization of $P^{(j)}_{post}(\phi)$ in order to maintain enough precision for each probability in the distribution (See Appendix B.4 for details). Lastly, the posterior distribution $P^{(j)}_{post}(\phi)$ for step j then becomes the prior distribution $P^{(j+1)}(\phi)$ for step $j+1$.

5.4 Results for Phase Estimation of Coherent States

We experimentally implement the measurement strategy for phase estimation of coherent states based on real-time optimization of the displacement operation for both objective functions with $L = 30$ adaptive steps and PNR(3). Figure 5.6 shows the experimentally measured variance calculated with Eq. (5.8) multiplied by the QFI for coherent states of $4|\alpha|^2$. The blue and orange points show the experimental variance for strategies which maximize the sharpness $\langle S(\beta, m)\rangle$ and mutual information $I(\beta, m)$, respectively. The points and error bars correspond to the average and standard deviation across five sets of $N = 10^4$ independent experiments with random initial relative phase $\phi_0 - \theta_0$ between the input and LO. We note that the experimental data is not corrected by the detection efficiency of our system. The solid purple and green lines show the expected variance obtained through Monte Carlo simulations of the phase estimation algorithm described above and including all experimental parameters. The red and black solid lines show the CRLB for ideal detection efficiency ($\eta = 1.0$) and our efficiency ($\eta = 0.70$), respectively. The red and black dashed lines show the ultimate limit of a heterodyne receiver: $1/\mathrm{CFI}_{het}$. We find that both strategies not only have almost the same performance, but both surpass the limit of an ideal

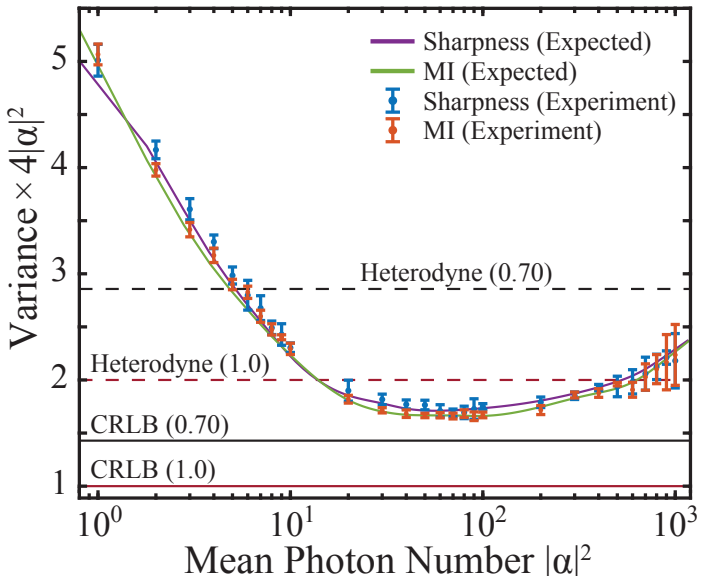

Figure 5.6: Experimentally obtained variance of the phase estimation measurements multiplied by $4|\alpha|^2$ (QFI). The points show the data for strategies maximizing the sharpness (blue) and mutual information (orange) and the error bars represent one standard deviation across five runs of $N = 10^4$ experiments each. We note that the data is not corrected for our overall efficiency of $\eta = 0.70$. The solid purple and green lines show the expected results for the sharpness and mutual information strategies given the experimental imperfections. The solid red and black lines show the CRLB without and with correction for our efficiency and the dashed lines show the limits of heterodyne detection. Figure adapted from: *Physical Review Letters* 125, 120505 (2020).

heterodyne measurement from $|\alpha|^2 = 20$ to $|\alpha|^2 \approx 600$ and an equivalent heterodyne measurement for $|\alpha|^2 \approx 6$ to $> 10^3$. The strategies reach uncertainties of less than 1.2 times the adjusted CRLB at $|\alpha|^2 \approx 70$ and less than 1.7 times the ideal CRLB. The increase in variance at larger input mean photon numbers is due to the limited number of adaptive steps L and PNR in our implementation, which can be avoided by either increasing the PNR or L [88].

We note that the increase in variance as the input mean photon number *decreases* is of a fundamental nature. This behavior is consistent with the "Mark II" adaptive homodyne strategy from Ref. [78] as well as the actual variance V_{can} for a so-called canonical measurement of the phase, implemented by projecting on to phase eigenstates $|\phi\rangle\langle\phi|$ [8, 78]. The variance V_{can} serves as a better limit on the performance of any

physical measurement compared to the CRLB, given that a canonical measurement is the optimal single-mode phase measurement [6, 78]. Furthermore, the CRLB serves only as a lower bound on the variance of any measurement strategy, and does not *guarantee* that a single-mode measurement reaches it [6]. Normalizing the measurement variance by V_{can} as opposed to the CRLB provides insight as to how different strategies perform compared to the best possible single-mode phase measurement [88].

Figure 5.7 shows the variance of both optimized non-Gaussian strategies in the ideal case ($\eta = 1.0, \xi = 1.0, \nu = 0$) for different numbers of adaptive steps and PNR but normalized by the variance of a canonical measurement V_{can}. The value of V_{can} at each input mean photon number is obtained using the framework in Ref. [78]. The dotted lines show both optimization strategies with $L = 30$ and PNR(3), the solid lines show $L = 200$ with PNR(3), and the dashed lines show $L = 30$ with PNR(12). The black line shows the performance of the best known adaptive homodyne strategy ("Mark II" from Ref. [78]) and the gray line shows the variance of an actual heterodyne measurement.

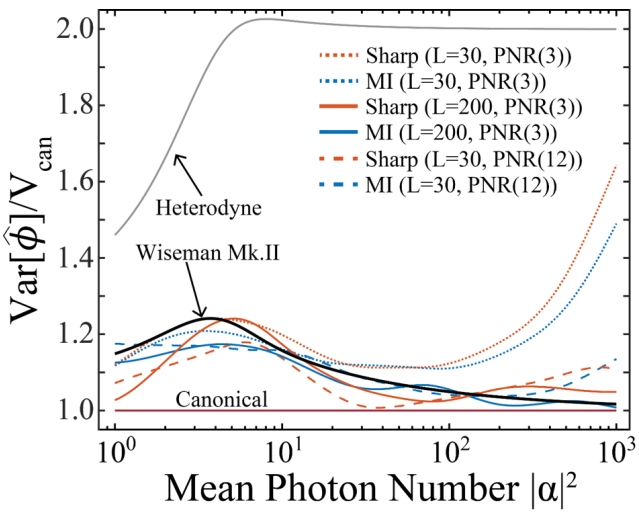

Figure 5.7: Variance of adaptive non-Gaussian strategies in the ideal case compared to that of a canonical phase measurement V_{can} (solid red). The orange and blue lines show the results for strategies which maximize the sharpness and mutual information, respectively, for three different values of L and PNR(m): $L = 30$ with PNR(3) (dotted), $L = 200$ with PNR(3) (solid), $L = 30$ with PNR(12) (dashed). The black solid line shows the variance of the "Mark II" adaptive homodyne strategy and the solid gray line shows the variance of an actual, ideal heterodyne measurement.

We find that in the ideal case, both optimization strategies provide similar performance to the "Mark II" adaptive homodyne receiver and the $L = 200$ with PNR(3) strategies have the best results at large input mean photon numbers. The strategies based on photon counting maintain similar performance to the "Mark II" scheme in the low mean photon number regime and are at most only 1.2 times the variance V_{can} of a canonical measurement at $|\alpha|^2 \approx 5$. Future work examining strategies based on photon counting with asymptotically many adaptive steps compared to homodyne based strategies will shed light on the limits of non-Gaussian vs. Gaussian approaches [88].

5.5 Gaussian Approximation and the Classical Fisher Information

We develop a Gaussian approximation to efficiently implement the real-time optimization of the displacement operation in the FPGA. The Gaussian approximation is based on calculating the variance of the full prior probability distribution $P^{(j)}(\phi)$ for adaptive step j and then approximating $P^{(j)}(\phi)$ as a Gaussian distribution:

$$\sigma^2 = \int_0^{2\pi} \phi^2 P^{(j)}(\phi) d\phi \;\rightarrow\; P^{(j)}(\phi) \approx P_G^{(j)}(\phi) = \mathcal{N} e^{-\frac{(\phi - \hat{\phi}_{MAP})^2}{2\sigma^2}}, \tag{5.9}$$

where $P_G^{(j)}(\phi)$ is the Gaussian approximation of $P^{(j)}(\phi)$. Figure 5.8 shows the optimal magnitude $|\beta|_{opt}$ and phase offset δ when using $P_G^{(j)}(\phi)$ in the function optimizations for different mean photon numbers $|\alpha|^2$ and variance σ^2. Panels (a) and (b) show the phase offset δ when optimizing mutual information and sharpness, respectively, and panels (c) and (d) show the optimal magnitude $|\beta|_{opt}$. The optimal values of the LO for both objective functions appear to asymptote to the same values as the prior variance σ^2 decreases.

We further investigate the asymptotic behavior of the optimal LO values by studying the classical Fisher information (CFI) for the measurement. The CFI characterizes

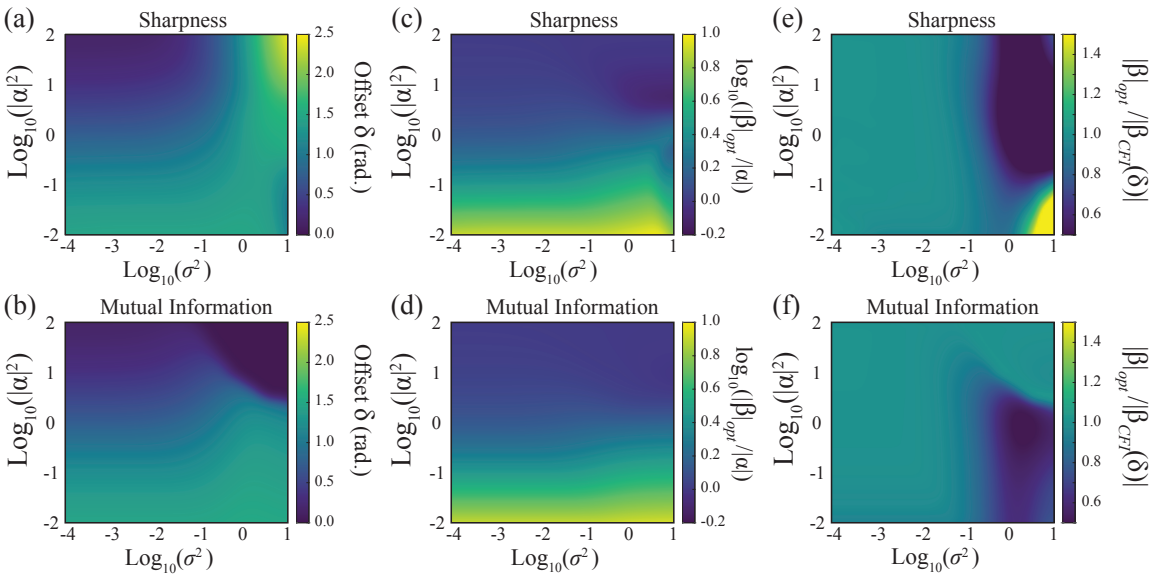

Figure 5.8: Optimal values from the Gaussian approximation for the (a,b) phase offset δ and (c,d) displacement magnitude $|\beta|_{opt}/|\alpha|$ for the sharpness (top panels) and mutual information (bottom panels) strategies. (e-f) show the ratio of the optimal value of $|\beta|_{opt}$ to the expected value which maximizes the Fisher information $|\beta_{CFI}(\delta)|$ in Eq. (5.11). The ratio asymptotes to a value of 1 as the prior distribution variance σ^2 decreases. Figure adapted from: *Physical Review Letters* 125, 120505 (2020).

how much information the the outcomes of a measurement contain about the unknown parameter ϕ, and can be viewed as a measure of how sensitive the likelihood functions of a measurement are to changes in the unknown parameter. The CFI for displaced photon counting $F(\phi)$ is a function of the unknown relative phase ϕ between the input state (ϕ_0) and LO (θ) and for infinite PNR is given by:

$$F(\phi) = \sum_{n=0}^{\infty} \mathcal{L}(n|\phi)\left(\frac{\partial}{\partial\phi}\ln\big(\mathcal{L}(n|\phi)\big)\right) = \frac{4|\alpha|^2|\beta|^2\sin^2(\phi)}{|\alpha|^2+|\beta|^2-|\alpha||\beta|\cos(\phi)}. \tag{5.10}$$

The CFI equals the quantum Fisher information (QFI) for coherent states when $\phi = 0$ and $|\beta| = |\alpha|$, corresponding to displacement to the vacuum state. However, our numerical investigations for adaptive non-Gaussian measurements reveal that even with hundreds of adaptive steps, a non-optimized strategy which attempts to displace the most likely state to the vacuum state cannot outperform a heterodyne measurement [87]. Nevertheless, optimizing the CFI with respect to the LO magnitude

reveals an optimality condition for $|\beta|$:

$$|\beta_{CFI}(\phi)|^2 = \frac{|\alpha|^2}{\cos^2(\phi)}, \tag{5.11}$$

and $F(\phi)$ when $|\beta| = |\beta_{CFI}(\phi)|$ is:

$$F(\phi)\Big|_{|\beta|=|\beta_{CFI}(\phi)|} = 4|\alpha|^2 = \text{QFI}. \tag{5.12}$$

Thus, the CFI for displaced photon counting is equal to the QFI as long as Eq. (5.11) is satisfied. In other words, for *any* relative phase ϕ between the input and LO, there is an optimal magnitude $|\beta_{CFI}(\phi)|$ which saturates the QFI.

One may expect the approximately optimal values for the LO to satisfy the optimality condition Eq. (5.11) when $j \gg 1$, given that $P_G^{(j)}(\phi)$ well describes $P^{(j)}(\phi)$ after many adaptive steps due to the small variance σ^2 and negligible skewness μ_3 of $P^{(j)}(\phi)$. The total relative phase between the unknown input (ϕ_0) and LO (θ_{opt}) in the limit of $P^{(j)}(\phi)$ having small variance is: $\lim_{\sigma\to 0} \phi_0 - \theta_{opt} \approx \Delta_{opt} \approx \delta$ due to $\hat{\phi} \approx \phi_0$ and $\mu_3 \approx 0$ such that $f(\mu_3) = 1$. The optimal displacement magnitude $|\beta|_{opt}$ and offset δ in the Gaussian approximation are obtained using $P_G^{(j)}(\phi)$ in the function optimizations. Therefore, we can obtain a prediction $|\beta_{CFI}(\delta)|$ of $|\beta|_{opt}$ based on δ and Eq. (5.11) as: $|\beta_{CFI}(\delta)|^2 = |\alpha|^2/\cos^2(\delta)$.

Figure 5.8(e,f) shows the ratio $|\beta|_{opt}/|\beta_{CFI}(\delta)|$ of the optimal $|\beta|_{opt}$ to the predicted $|\beta_{CFI}(\delta)|$ for strategies optimizing the sharpness and mutual information, respectively. We find that the predicted value $|\beta_{CFI}(\delta)|$ asymptotes to the optimal value $|\beta|_{opt}$ for $\log_{10}(\sigma^2) < -1.5$ for all mean photon numbers for both objective functions. We note that the optimality condition in Eq. (5.11) is obtained assuming infinite PNR, but the $|\beta|_{opt}$, δ, and $|\beta_{CFI}(\delta)|$ are all obtained assuming a PNR(3) detector, as in the experiment. This result that $\lim_{\sigma\to 0} |\beta_{CFI}(\delta)| = |\beta|_{opt}$ validates the Gaussian approximation and also shows that the LO values obtained through this approximation are ones which saturate Eq. (5.11) and allow the CFI for displaced photon counting to equal the QFI for coherent states.

We also study the function landscape for both the sharpness and mutual information, as opposed to just their optimal values in the Gaussian approximation.

Figure 5.9(a) shows an example prior probability distribution $P^{(j)}(\phi)$ for an adaptive step early ($j = 4$) in the measurement, which in general will be relatively broad and not well approximated by a Gaussian distribution. The blue dashed line shows the true value of the phase ($\phi_0 = \pi$) and the red line shows the current phase estimate $\hat{\phi}$ from Eq. (5.7). The red and black lines show the optimal LO phase values $\theta_{opt,S}$ and $\theta_{opt,MI}$ which maximize the sharpness and mutual information, respectively. The optimal LO values for the sharpness (Eq. (5.2)) or mutual information (Eq. (5.5)) correspond to global maxima in the function landscapes and are solely a function of the full prior probability distribution $P^{(j)}(\phi)$.

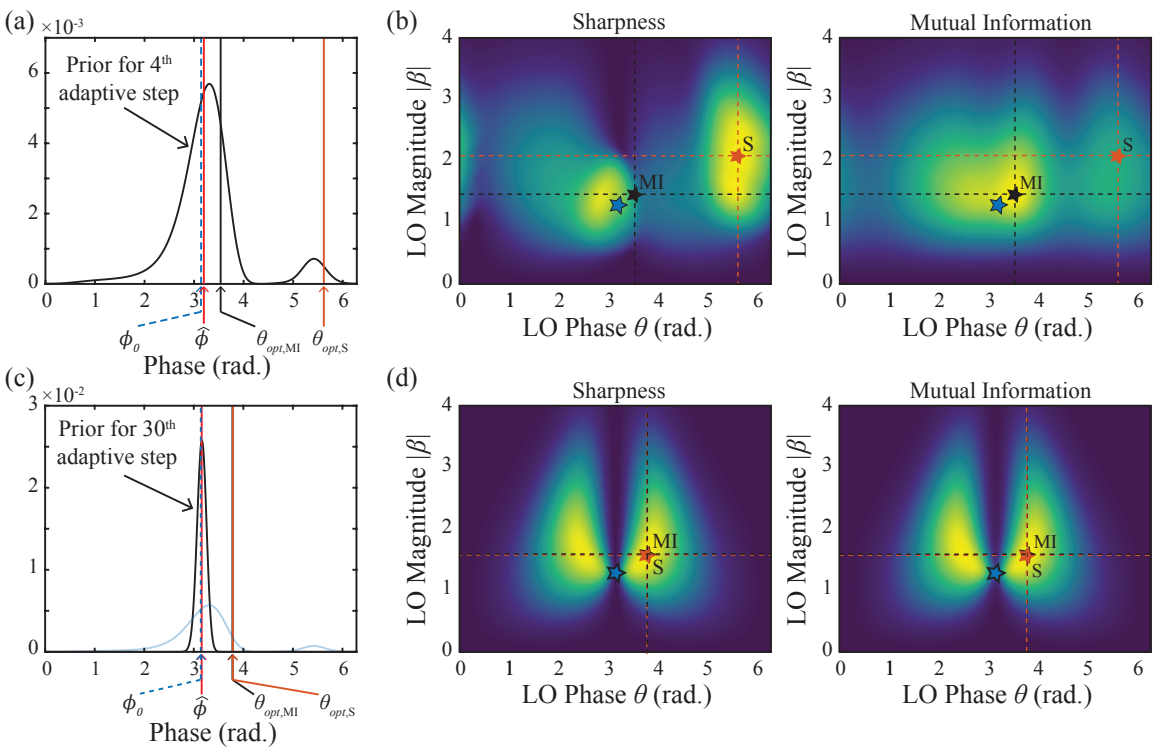

Figure 5.9: Comparison between the sharpness and mutual information function landscapes and optimal values $|\beta|_{opt}$ and θ_{opt} for different prior probability distributions. (a) A prior distribution for an early adaptive step ($j = 4$) will generally be non-Gaussian and (b) the function landscapes and optimal LO values (stars) will be very different. (c) A prior distribution for the last adaptive step ($j = L = 30$) will be well approximated by a narrow Gaussian distribution. (d) In this limit of $\sigma \to 0$ the entire function landscape and optimal values are almost identical for both objective functions.

Figure 5.9(b) shows the function landscapes for the sharpness (left panel) and mutual information (right panel) given the prior probability distribution in Fig. 5.9(a), and are calculated assuming an ideal experiment with PNR(3). The orange and black stars show the optimal LO values ($|\beta|_{opt}$ and $\theta_{opt} = \hat{\phi} + \Delta_{opt}$) which maximize the sharpness (S) and mutual information (MI), respectively, and the blue point corresponds to displacing the current hypothesis to the vacuum state ($|\beta| = |\alpha|, \theta = \hat{\phi}$). It is clear that the two optimal values and the function landscapes are very different from one another. This difference is due to the non-Gaussian form of the prior probability distribution combined with the fact that the sharpness and mutual information are simply very different functions.

On the other hand, Figure 5.9(c) shows an example of a prior probability distribution for the last adaptive step ($j = L = 30$) of the measurement, which is very well approximated as a Gaussian distribution with small variance ($\sigma \to 0$). The optimal LO phase values $\theta_{opt,S}$ and $\theta_{opt,MI}$ are now almost identical to each other, as predicted based on Fig. 5.8. Figure 5.9(d) shows the function landscapes for the sharpness (left panel) and mutual information (right panel) using the prior probability distribution in Fig. 5.9(c). We find that the globally optimal LO values ($|\beta|_{opt}$ and $\theta_{opt} = \hat{\phi} + \Delta_{opt}$) which maximize each objective function are almost identical, which is consistent with the Gaussian approximation results discussed above. Furthermore, it is apparent that the *entire* function landscape for both object functions are almost identical. This similarity is surprising given how different the two objective functions are and suggests that the sharpness and mutual information may be more similar than they appear in the limit of $\sigma \to 0$, but requires further investigation [88].

Our studies showing that both strategies have not only the same optimal values but also the same optimization landscapes in the limit of a narrow Gaussian prior distribution is consistent with the theoretical work in Ref. [246]. Reference [246] shows that Bayesian experimental designs which maximize the mutual information, or posterior variance (i.e. sharpness), are asymptotically efficient and can saturate the CRLB [250]. In addition, we note that the Kullback-Liebler divergence in Eq.

(5.6) has a series expansion in the limit of the two distributions being close to one another, and the first non-zero term in that expansion is in fact the classical Fisher information [248, 251]. Therefore, maximizing the mutual information is analogous to maximizing the Fisher information in the limit of $\sigma \to 0$, but with more information contained in the higher order expansion terms. On the other hand, maximizing the expected sharpness will achieve minimal Holevo variance [249]. An interpretation of this result is that each strategy (sharpness and mutual information) attempts to reach the CRLB but through two different approaches of either maximizing the Fisher information, or minimizing the variance [88].

5.6 Conclusion

We develop and demonstrate a strategy for single-shot phase estimation of optical coherent states which surpasses the limit of ideal heterodyne detection [87]. This measurement is based on adaptive displaced photon counting and real-time optimization of the displacement operation conditioned on the measurement history. We study two different objective functions for the optimization, the sharpness and mutual information, and find that they produce almost identical results in terms of the overall measurement variance. We also develop and analyze a Gaussian approximation which allows the experimental implementation of the strategy to obtain *approximately* optimal LO values based on just the variance and skewness of the prior probability distribution in real-time for each adaptive step. This approach as well as a quasi-normalization procedure within the Bayesian updating allows the FPGA to update the distribution for the input phase and obtain the optimal displacement values in only $\approx 1\ \mu s$. We implement this measurement across a wide range of input energies for a strategy with $L = 30$ adaptive steps and PNR(3). Our demonstration achieves variances of less than 1.2 times the adjusted CRLB and outperforms the ideal heterodyne limit from $|\alpha|^2 \approx 20$ to ≈ 600. Furthermore, we analyze the performance of ideal strategies with large numbers of adaptive steps and

high photon number resolution. We find that these strategies can perform at least as well as the best known adaptive homodyne measurements and reach variances close to that of the best possible physical measurement.

Chapter 6

Maximizing Mutual Information with Higher Order Encodings

A significant result in information theory is that communication strategies based coherent states of light can reach the ultimate limits of information transfer, known as the capacity, for a pure loss channel [18, 71, 252–254]. The capacity of a particular channel is given by the maximum possible mutual information optimized over all physical states, encodings, and POVMs, under an average energy constraint, and also referred to as the Holevo bound [153, 253]. To reach the capacity for a pure loss channel, a sender needs to encode information into optimal code-words created from coherent states with non-uniform probability [18]. However, the optimal POVM for decoding the information requires joint measurements over the elements of the codewords at the receiver [20, 255], a task which has yet to be fully experimentally realized [256–258]. Experiments have demonstrated strategies based on code-word nulling [259] but such strategies still require incorporation of quantum non-demolition measurements to approach the capacity for a pure loss channel. On the other hand, the limit of heterodyne receivers and coherent states with Gaussian modulation in a symbol-by-symbol paradigm is known as the Shannon limit [5, 20, 252, 260, 261], and serves as a benchmark for the limits of conventional technology. Studying the role of

non-Gaussian measurements in maximizing the amount of information transfer in this symbol-by-symbol approach will yield insights into how to surpass the Shannon Limit using quantum measurements.

In a traditional communication scenario, a sender encodes information into a finite set of M coherent states and thus the maximum amount of mutual information attainable is $\log_2(M)$. A strategy which utilizes $M = 16$ states can encode four bits of information by using both the phase and amplitude of the coherent states, as discussed in Section 1.3. Furthermore, the mutual information can be increased by imposing a Gaussian prior probability distribution onto the states [18, 261]. However, the amount of information transfer possible when using heterodyne measurements is bounded by the Shannon limit. Recent theoretical work in Ref. [262] has shown that adaptive non-Gaussian receivers for a 16-QAM alphabet can surpass the limits of ideal heterodyne detection and Gaussian modulation. This performance is achieved by optimizing the prior probability distribution for the states combined with adaptive hypothesis testing based on displacement operations and photon counting [51, 117].

In this chapter, we extend the analysis presented in Ref. [262] to include both experimental imperfections as well as photon number resolution. We find that receivers based on displacement operations, photon counting, and feedback using a 16-QAM alphabet can surpass the limits of equivalent heterodyne detection with a feasible number of adaptive measurement steps with PNR(1). Furthermore, a receiver with photon number resolving capabilities requires a smaller number of adaptive steps to reach the same performance. We also discuss the potential experimental implementation of these measurement strategies.

6.1 Measurement Strategy

Reference [262] shows that communication strategies based on coherent states using an optimized prior probability distribution and adaptive non-Gaussian measurements

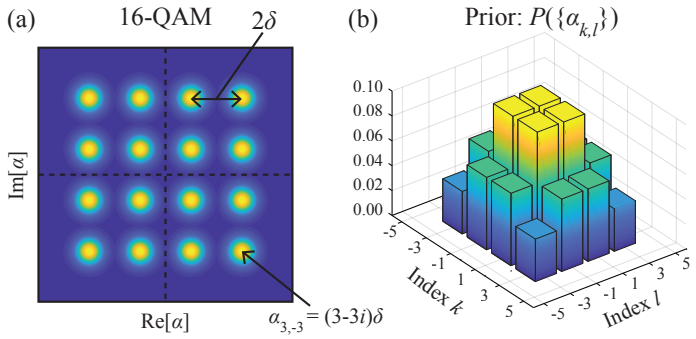

Figure 6.1: (a) 16-QAM encoded coherent states $|\alpha_{k,l}\rangle$ in phase space where δ controls the spacing of the alphabet. (b) Example Gaussian prior probability distribution for the 16-QAM alphabet in (a).

can outperform the limits of heterodyne detection and Gaussian modulation. We focus on a scenario where the sender and receiver use a 16-QAM alphabet, which is a 4×4 grid of coherent states: $|\alpha_{k,l}\rangle = |(k+il)\delta\rangle$ with $k,l \in \{-3,-1,1,3\}$. Figure 6.1(a) shows the 16-QAM alphabet $|\alpha_{k,l}\rangle$ in phase space. The parameter δ is the spacing of the states in phase space and controls the energy of the alphabet. In contrast to other state discrimination strategies discussed in this thesis, strategies for maximizing information transfer use an input alphabet with the prior probabilities determined by a Gaussian distribution: $P(\alpha_{k,l}) = \mathcal{N} e^{-|\alpha_{k,l}|^2/(2\sigma^2)}$ [18, 262]. Figure 6.1(b) shows an example Gaussian prior probability distribution applied to the 16-QAM alphabet. Here, σ and δ are optimized to achieve maximal mutual information for a given average energy and for a specific measurement strategy.

Figure 6.2 shows the experimental concept for maximizing information transfer using non-Gaussian measurements. A sender (Alice) first encodes information into a 16-QAM alphabet $\{|\alpha_{k,l}\rangle\}$ with a Gaussian prior probability distribution $P(\{\alpha_{k,l}\})$ (left 3D plot). The coherent states are then transmitted through a lossy communication channel and measured by an adaptive non-Gaussian receiver based on displacement operations, single photon detection, and feedback. The receiver implements an adaptive Bayesian strategy with L adaptive steps and PNR(m), where the magnitude $|\beta| \in \{|\alpha_{k,l}|\}$ and phase $\arg(\beta) \in \{\arg(\alpha_{k,l})\}$ of the displacement $\hat{D}(\beta)$ are set to displace the most likely state to the vacuum state, followed by photon counting [117].

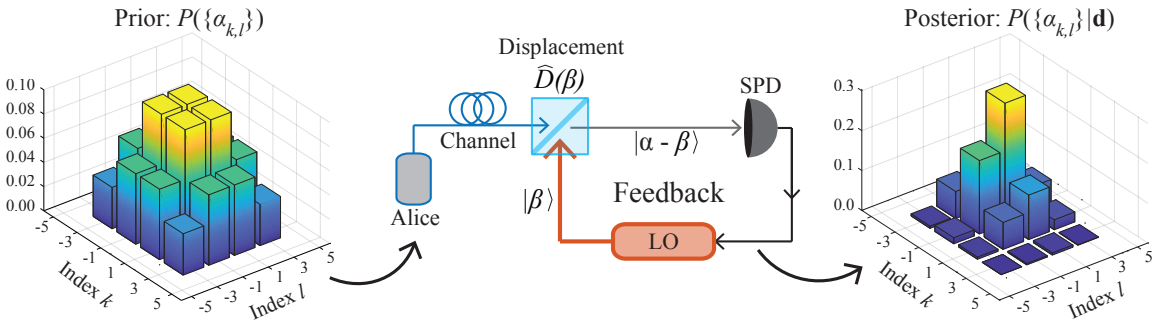

Figure 6.2: A sender (Alice) encodes information into a 16-QAM alphabet $\{|\alpha_{k,l}\rangle\}$ with a Gaussian prior probability distribution $P(\{\alpha_{k,l}\})$ (left 3D plot). The transmitted state is measured by an adaptive non-Gaussian receiver where $P(\{\alpha_{k,l}\})$ is incorporated into the Bayesian updating strategy as the initial prior distribution. The entropy of the posterior distribution $P(\{\alpha_{k,l}|\mathbf{d}\})$ (right 3D plot) after measurement is smaller than that of the prior distribution. The mutual information is equal to the average difference between the entropy of the prior and posterior probability distributions, averaged over all possible detections.

The final posterior distribution $P(\{\alpha_{k,l}\}|\mathbf{d})$ (right 3D plot) given a particular detection history $\mathbf{d} = \{d_1, ..., d_L\}$ for the measurement will have an entropy smaller than that of $P(\{\alpha_{k,l}\})$. The receiver strategy aims to achieve the maximum mutual information $I(\bar{n}, \sigma, \delta)$ between the sender and receiver for a given average energy $\bar{n}$, where $I(\bar{n}, \sigma, \delta)$ is obtained as:

$$I(\bar{n}, \sigma, \delta) = \sum_{\{\mathbf{d}\}} \sum_{k,l=1}^{4} P(\alpha_{k,l}, \mathbf{d}) \log_2 \left[\frac{P(\alpha_{k,l}, \mathbf{d})}{P(\alpha_{k,l}) P(\mathbf{d})} \right], \tag{6.1}$$

$$= \sum_{\{\mathbf{d}\}} \sum_{k,l=1}^{4} P(\alpha_{k,l}|\mathbf{d}) P(\mathbf{d}) \log_2 \left[\frac{P(\alpha_{k,l}|\mathbf{d})}{P(\alpha_{k,l})} \right], \tag{6.2}$$

$$= \mathrm{H}\big[P(\{\alpha_{k,l}\})\big] - \sum_{\{\mathbf{d}\}} P(\mathbf{d}) \mathrm{H}\big[P(\{\alpha_{k,l}\}|\mathbf{d})\big]. \tag{6.3}$$

Here $P(\alpha_{k,l}, \mathbf{d}) = P(\alpha_{k,l}|\mathbf{d})P(\mathbf{d})$ is the joint probability for the state $|\alpha_{k,l}\rangle$ and detection $\mathbf{d}$. The summation over $\{\mathbf{d}\}$ is over all possible detection histories, the summation over k and l is over all possible input states, and

$$\mathrm{H}[P(\{\alpha_{k,l}\})] = \sum_{k,l} P(\alpha_{k,l}) \log_2 \big[P(\alpha_{k,l})\big], \tag{6.4}$$

$$\mathrm{H}[P(\{\alpha_{k,l}\}|\mathbf{d})] = \sum_{k,l} P(\alpha_{k,l}|\mathbf{d}) \log_2 \big[P(\alpha_{k,l}|\mathbf{d})\big], \tag{6.5}$$

are the entropy of the prior $P(\{\alpha_{k,l}\})$ and posterior distribution $P(\{\alpha_{k,l}\}|\mathbf{d})$, respectively. The average energy $\bar{n}$ of the coherent state alphabet is given by:

$$\bar{n} = \sum_{k,l} P(\alpha_{k,l})|\alpha_{k,l}|^2 = \sum_{k,l} \frac{e^{-\delta^2(k^2+l^2)/(2\sigma^2)}}{\sum_{m,n} e^{-\delta^2(m^2+n^2)/(2\sigma^2)}} \times \delta^2(k^2+l^2), \tag{6.6}$$

which depends solely on σ and δ. Note that by fixing the value of $\bar{n}$ and σ, the value of δ is obtained as the value which satisfies Eq. (6.6).

There is an optimal value of $\sigma = \sigma_{opt}(\bar{n})$ for each value of the average energy $\bar{n}$ that maximizes the mutual information $I(\bar{n}, \sigma, \delta)$ given that the receiver implements the adaptive non-Gaussian measurement described above. We numerically find $\sigma_{opt}(\bar{n})$ using Monte Carlo simulations of the non-Gaussian measurement. First, we choose a particular value of $\bar{n}$ and then scan over a range of values for σ. We obtain the corresponding value of δ by inverting Eq. (6.6) for every $\bar{n}$ and σ, thus fixing $\{|\alpha_{k,l}\rangle\}$ and $P(\{\alpha_{k,l}\})$. We then simulate 5×10^4 independent measurements and calculate the mutual information $I(\bar{n}, \sigma, \delta)$ for these values of $\bar{n}$, σ, and δ using Eq. (6.3).

Figure 6.3(a) shows the mutual information as a function of σ for a fixed average

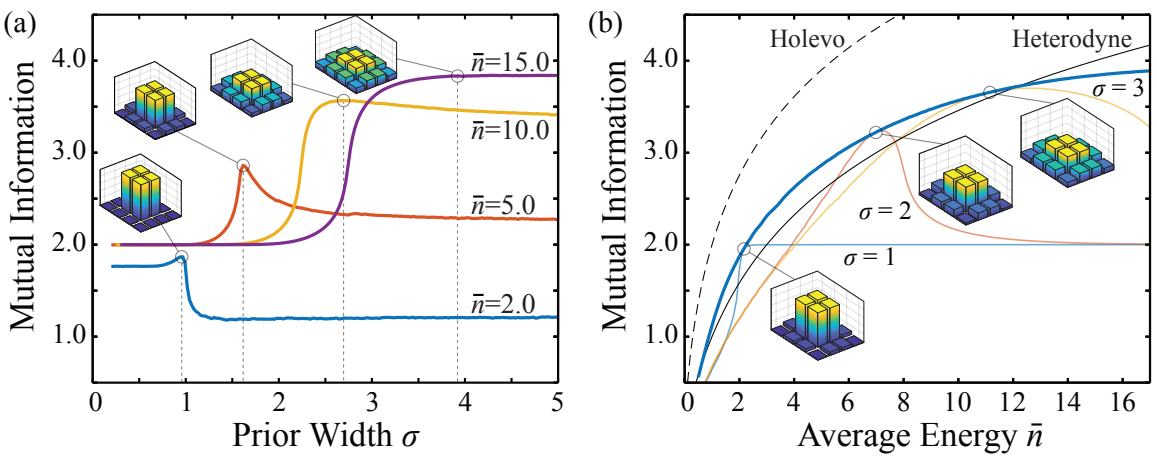

Figure 6.3: (a) Mutual information for an ideal receiver with $L = 16$ and PNR(1) as a function of σ for $\bar{n}$=2, 5, 10, and 15 in blue, orange, yellow, and purple, respectively. The vertical dashed lines show the value of $\sigma_{opt}(\bar{n})$ for each $\bar{n}$ and the 3D insets show $P(\{\alpha_{k,l}\})$ for these values. (b) The dark blue line shows the maximal mutual information $I(\bar{n}, \sigma_{opt}(\bar{n}), \delta_{opt}(\bar{n}))$ as a function of the average energy. The black and gray lines show the limit of heterodyne detection with coherent states and Holevo bound, respectively.

energy $\bar{n}$=2, 5, 10, and 15, for a ideal strategy with $L = 16$ adaptive steps and PNR(1). We extract the value of $\sigma_{opt}(\bar{n})$ from these curves for each $\bar{n}$ as the value of σ which achieves maximum mutual information. The 3D panels show the prior probabilities $P(\{\alpha_{k,l}\})$ for the states at $\sigma_{opt}(\bar{n})$ for each energy $\bar{n}$. Figure 6.3(b) shows the maximum achievable mutual information $I(\bar{n}, \sigma_{opt}(\bar{n}), \delta_{opt}(\bar{n}))$ as a function of $\bar{n}$ for an ideal non-Gaussian strategy based on adaptive photon counting with PNR(1) and $L = 16$ (blue). The thin solid lines show the mutual information as a function of $\bar{n}$ for fixed values of σ =1,2, and 3, in blue, orange, and yellow, respectively. The 3D plots show the optimal prior probability distribution at the average energy $\bar{n}$ for which $\sigma_{opt}(\bar{n})$=1, 2, and 3. The black solid line shows the limit for heterodyne measurements with a Gaussian modulation [261,262]. The black dashed line shows the Holevo bound, which is the ultimate limit of information transfer and requires optimal code-words and joint measurements at the receiver [6, 18, 153].

6.2 Proposed Experimental Implementation

Figure 6.4 shows a diagram of the proposed experimental setup for the 16-QAM state discrimination measurement we aim to demonstrate in the future. We will include an amplitude modulator in the signal arm of the interferometer and both the phase (PM) and amplitude (AM) modulators for the signal will each be controlled by an 8-bit digital-to-analog converter (DAC). This setup will allow for easily changing the magnitude as well as phase of the signal with the field-programmable gate array (FPGA) to prepare different input states in the 16-QAM constellation. We will enforce the prior probabilities $P(\{\alpha_{k,l}\})$ by constructing a look-up table (LUT) which sequentially prepares the states $|\alpha_{k,l}\rangle$ in the signal arm for $N_{k,l}$ independent experiments each, and we will obtain data in blocks of $\sum_{k,l} N_{k,l} = 2^{11} = 2048$ experiments. The signal LUT will loop over the states $|\alpha_{k,l}\rangle$ such that $|\alpha_{1,1}\rangle$ is prepared for the first $N_{1,1}$ experiments, followed by $|\alpha_{1,2}\rangle$ for $N_{1,2}$ experiments, and so on, and this pattern will repeat every 2^{11} experiments. The number of experiments $N_{k,l}$ per input state is

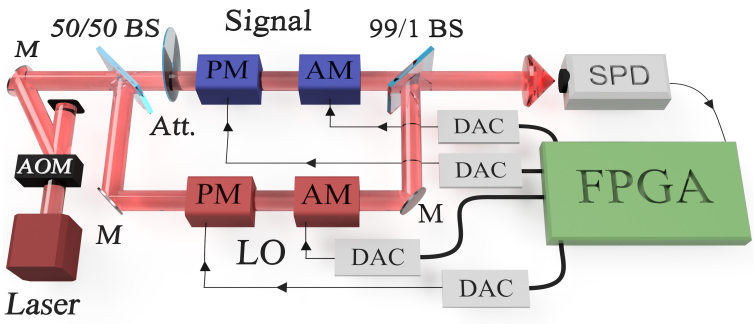

Figure 6.4: Proposed experimental setup for maximizing mutual information with a 16-QAM alphabet. The signal and LO phase and amplitude will each be controlled by a modulator connected to an 8-bit digital-to-analog converter (DAC). The field-programmable gate array (FPGA) will implement the hypothesis testing strategy and prepare the input state using the phase and amplitude modulators.

obtained as: $N_{k,l} = \text{round}\big(2^{11} P(\alpha_{k,l})\big)$, where round$(\cdot)$ rounds to the nearest integer such that $N_{k,l} / \sum_{m,n} N_{m,n} \approx P(\alpha_{k,l})$. We will prepare the displacement $\hat{D}(\beta)$ in a similar way to previous experiments where a second pre-calculated LUT stores the amplitude $|\beta|$ and phase $\arg(\beta)$ for each adaptive step as discussed below. We will construct this second LUT to implement the hypothesis testing approach where the receiver attempts to displace the current hypothesis to the vacuum state.

6.3 Adaptive Steps vs. PNR trade-off

There exists a practical trade-off between photon number resolution and number of adaptive steps in an experimental demonstration. Implementing these measurement strategies at high bandwidth requires using a LUT within the FPGA to store the displacement value β given every possible detection history $\mathbf{d}$ [51, 83, 117]. Each adaptive step requires an individual LUT, and the table for adaptive step j requires $(m+1)^{j-1}$ elements for a strategy with PNR(m). For example, $j = 2$ requires $m+1$ elements because there are $m+1$ possible photon detection results from the first

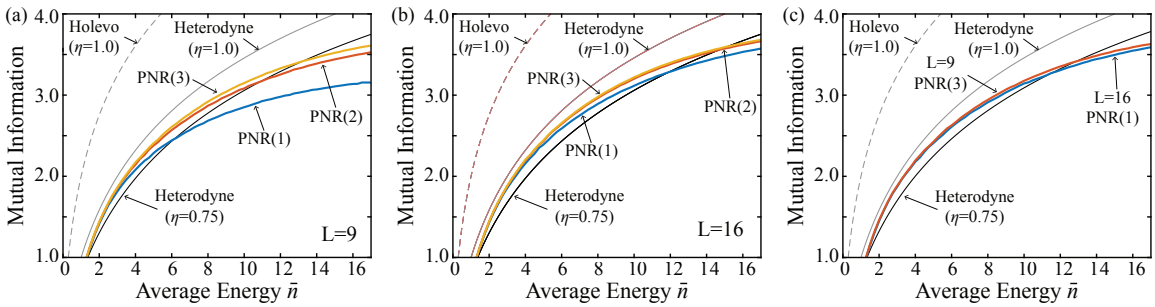

Figure 6.5: Comparison between different non-Gaussian strategies. (a) The blue, orange, and yellow lines show the mutual information for a strategy with $L = 9$ adaptive steps and PNR(1), PNR(2), and PNR(3), respectively. (b) The colored lines show the performance of a strategy with $L = 16$ adaptive steps for each PNR. (c) Comparison between the $L = 9$ with PNR(3) and $L = 16$ with PNR(1) strategies from (a) and (b). A strategy with less adaptive steps but larger PNR allows for slightly higher mutual information.

adaptive step, and the third adaptive step $j = 3$ requires $(m + 1)^2$ elements because there are $m + 1$ outcomes each for $j = 1$ and $j = 2$. Note, the initial displacement values for the first adaptive step $j = 1$ are pre-determined and do not require a LUT. Therefore, the *total* size (in bits) of the LUTs required to store the displacement magnitude and phase when they are each represented by eight bit numbers is:

$$2 \times 8 \times \sum_{j=2}^{L} (m+1)^{j-1} = 16 \times \frac{(m+1)^L - (m+1)}{m}. \tag{6.7}$$

Constructing a LUT in the FPGA requires using a certain amount of internal memory, which comes in specified block sizes and there is a limited total amount available. Thus, there is a practical question of whether to use a high L but low PNR strategy, or the other way around. This trade-off exists because increasing the PNR will decrease the number of possible adaptive steps, and vice versa, due to the finite amount of memory available. A strategy with PNR(1) can be implemented with at most $L = 16$ ($16 \times (2^{16} - 2) \approx 1.05$ Mb) adaptive steps, and a PNR(3) strategy can be implemented with at most $L = 9$ adaptive steps ($16 \times (4^9 - 4)/3 \approx 1.40$ Mb) given the total amount of memory in the FPGA in our setup (≈ 2 Mb).

Figure 6.5(a-c) shows the trade-off in performance between different values of L and PNR in a realistic experiment with detection efficiency $\eta = 0.75$, visibility

$\xi = 0.997$, and dark count rate $\nu = 160/\text{s}$. Figure 6.5(a) shows the mutual information for $L = 9$ with PNR(m) for m=1, 2, and 3, in blue, orange, and yellow, respectively and 6.5(b) shows the mutual information for $L = 16$ with different PNR. Figure 6.5(c) shows the comparison between the $L = 9$ with PNR(3) strategy (blue) from (a) and the $L = 16$ with PNR(1) strategy from (b). We find that it is more beneficial to implement a smaller number of adaptive steps with a larger PNR, which achieves higher mutual information as the average energy increases.

6.4 Future Directions

Experimentally demonstrating the measurement strategy discussed above will show that non-Gaussian receivers based on adaptive photon counting can surpass the limits of conventional measurements in terms of information transfer. Furthermore, communication strategies which maximize mutual information may also benefit from optimizing the displacement operations using global optimizations [83] or reinforcement learning approaches, both described in Ch. 3. Global optimizations of the measurement strategy will aim to achieve the highest possible amount of mutual information, but may become impractical for large numbers of adaptive steps with PNR capabilities. The reinforcement learning framework would allow for finding optimized measurement strategies with a large number of adaptive steps and moderate PNR. In addition, the prior probability distribution of the input states may be incorporated into the reinforcement learning optimization. This inclusion would allow the receiver to find optimized communication strategies for channels with unknown noise properties, or when conventional optimization is impractical or even impossible.

6.5 Conclusion

In this chapter, we extend the theoretical work in Ref. [262], where the authors investigate adaptive non-Gaussian measurements for maximizing mutual information. Reference [262] shows that non-Gaussian measurements can outperform strategies based on heterodyne detection and Gaussian modulation. Our numerical simulations show that adaptive non-Gaussian receivers combined with a 16-QAM alphabet and an optimized prior probability distribution can potentially surpass the limits of equivalent heterodyne measurements when accounting for experimental imperfections. In addition, non-Gaussian receivers which employ photon number resolution can achieve the same performance with fewer adaptive measurement steps. Specifically, a strategy with $L = 9$ adaptive steps and PNR(3) can outperform a strategy with $L = 16$ and PNR(1), based on our analysis including experimental limitations. In the future, we plan to experimentally demonstrate these strategies which maximize the mutual information using adaptive non-Gaussian measurements. Furthermore, machine learning techniques discussed in Chapter 3 can potentially be applied to adaptive non-Gaussian receivers to find optimized strategies in parameter regimes where conventional approaches are impractical, or even for channels with unknown noise properties.

Chapter 7

Thesis Conclusion

Non-Gaussian measurements can surpass the sensitivity limits of conventional technologies based on homodyne and heterodyne detection. In this thesis, we show that the non-Gaussian element corresponding to single-photon detection provides a benefit across a wide range of state discrimination and parameter estimation tasks. Quantum measurements based on photon counting can also follow an adaptive approach. Instead of measuring the entire input state at once, a measurement strategy can perform many different measurements by splitting the input state into many temporal modes and sequentially measuring each mode. The information gained by each sequential detection is maintained through recursive Bayesian updating of the probability for each possible input state given the entire detection history. Furthermore, non-Gaussian strategies can be optimized using the prior information about the input state or in real-time. This possibility of optimized measurements can enhance the performance of different strategies for various tasks. We experimentally and theoretically show that non-Gaussian measurements can be used for diverse problems and can to outperform conventional strategies in terms of discrimination error, mutual information, and estimation variance.

We investigated non-adaptive measurements for binary phase-shift-keyed (BPSK) coherent states using an optimized Kennedy receiver combined with photon number

resolution (PNR). We found that PNR allows the receiver to maintain sub-QNL performance while providing robustness to experimental noise and imperfections, such as non-ideal interference visibility and detector dark counts. We also applied this optimized Kennedy receiver with PNR to the problem of communication over channels with complex noise properties, in particular that of phase diffusion. While communication strategies based on a BPSK encoding and non-Gaussian measurements provide some benefit over the QNL when the noise strength of the channel is small, the performance of such strategies is quickly degraded as the noise level increases. To this end, we investigated an optimized communication strategy, where both the input coherent state alphabet as well as the displacement operation are jointly optimized and tailored to provide the best performance for a particular noise level. The optimized communication strategy allows the receiver to recover sub-QNL performance even in the presence of severe channel noise. We also implemented a more general class of binary measurements, which goes beyond the familiar paradigm of minimum error state discrimination. These more general measurements achieve the smallest possible error probability for a fixed probability an an inconclusive result. We experimentally validated this measurement for binary coherent states, which is implemented using displacement operations, photon counting, and fast feedback, where the displacement amplitude follows a particular optimal waveform.

Increasing the alphabet size to four coherent states increases the total amount of information which can possibly be transmitted. We studied three different measurement strategies for the discrimination of quaternary phase-shift-keyed (QPSK) coherent states: (1) simultaneous hypothesis testing, (2) global optimizations for adaptive approaches, and (3) reinforcement learning of an alternative receiver architecture. On one hand, "minimum resource" measurements which are non-adaptive, have no PNR capabilities, and use the smallest number of simultaneous detections potentially enable measurement strategies for sub-QNL discrimination at high bandwidth. On the other hand, global optimizations allow for adaptive strategies to outperform the QNL in the single photon regime compared to non-optimized strategies. As an

alternative approach, machine learning provides a path for constructing novel receivers for state discrimination based on artificial neural networks which are recursive and are able to maintain an internal memory of the detection history. We show that these networks can be trained to replace simple state discrimination measurements with no reduction in performance. Furthermore, reinforcement learning can be used to find optimized measurement strategies in situations where conventional techniques for multi-parameter optimization are impractical or impossible. As a further investigation for non-Gaussian measurements for multiple coherent states for communications, we investigated scenarios where a sender and receiver can potentially transmit four bits of information utilizing a 16-QAM alphabet. This alphabet employs the phase and amplitude of the coherent states simultaneously to encode information, as opposed to just the phase as with BPSK and QPSK encodings. We extend previous theoretical work and show that adaptive non-Gaussian measurements can potentially surpass the limit of heterodyne-based strategies. This performance is achieved by using a prior probability distribution which is optimized to maximize the mutual information under an average every constraint.

Noise in a communication channel can severely degrade the ability of a receiver to perform state discrimination, and reduce information transfer in communications. Estimating and tracking of channel noise is a crucial task for any communication scenario and advanced techniques currently facilitate communication with conventional measurement strategies. However, these standard approaches cannot be applied to adaptive non-Gaussian measurements. We developed an estimation algorithm to implement real-time phase tracking in adaptive non-Gaussian receivers for QPSK coherent states. We experimentally demonstrated that an intuitive estimator enables sub-QNL state discrimination in the presence of dynamic phase noise. In addition, we developed a method for tracking multiple channel noise parameters based on a neural network estimator. We numerically showed that a properly trained neural network allows for sub-QNL performance in the presence of both phase and amplitude noise across a wide range of channel noise levels. This approach is far more computationally

efficient than conventional estimation techniques and can potentially enable high-bandwidth discrimination beyond the limits of heterodyne detection in realistic communication channels.

Estimation of the phase of a coherent state of light is a central problem in metrology, and the best possible (canonical) measurement strategy has no known physical implementation. We showed that adaptive non-Gaussian measurements can outperform the limits of heterodyne detection and approach the ultimate bounds on the uncertainty in estimation of the phase of a coherent state. These adaptive strategies are based on real-time optimization of the displacement operation, photon number resolution, and fast feedback. We demonstrated two optimization strategies and develop an optimization algorithm which operates in real-time based on a Gaussian approximation. We found that both optimization strategies provide uncorrected estimation variances below the limit of ideal heterodyne detection, and that the two strategies have similar behavior in the Gaussian approximation.

The work presented in this thesis serves as a foundation for addressing the seemingly simple question of how to implement a *good* measurement. The answer however, is quite complex and depends on how we define *good*. On one hand, quantum measurements which discriminate between finite sets of coherent states better than heterodyne detection can enable rates on information transfer which are inaccessible to conventional technologies. On the other hand, the novel techniques for combating channel noise discussed here provide insight on how to readily transfer optimized photon counting measurements for coherent states out of the lab. One particular open question is what truly are the limits of non-Gaussian measurements which are optimized in real-time, and the algorithms and machine learning techniques explored in this thesis may help shed light on these limits. So how does one actually implement a good, or even the best possible measurement? It's certainly challenging, but an *excellent* place to start is with counting photons.

Chapter 8

Appendices

Appendix A

Adaptive Photon Counting

A.1 Measurement Strategy

The non-Gaussian element of photon counting can enable non-conventional measurement strategies to surpass the QNL. However, the performance of measurements based on photon counting can be further improved by allowing the receiver to implement an adaptive strategy. Instead of measuring the entire input pulse all at once, the receiver splits the input state into L temporal modes and measures each mode sequentially. Here, measurement of a single temporal mode is referred to as a single adaptive step, each of which is comprised of a coherent displacement operation $\hat{D}(\beta)$ followed by photon counting. The receiver must first choose the displacement parameter β at the beginning of each adaptive step j of L total steps. The prior probability distribution $P(\{\alpha_k\})$ for the set of coherent states $\{\alpha_k\}$ at step j can be used to infer the best displacement operation $\hat{D}(\beta)$. A typical approach in state discrimination is to implement a hypothesis testing procedure, where β is chosen such that current hypothesis for the input state is displaced to the vacuum state, i.e. $\beta = \text{argmax}_{\alpha_k} P(\{\alpha_k\})$. The measurement then detects a particular number of photons n from the displaced input state $\hat{D}(\beta)|\alpha_k\rangle = |\alpha_k - \beta\rangle$. Then, the receiver calculates the posterior distribution $P(\{\alpha_k\}|n, \beta)$ given the measurement

result of n detected photons during adaptive step j through Bayes rule:

$$P(\{\alpha_k\}|n,\beta) = \frac{1}{P(n)}\mathcal{L}(n|\{\alpha_k\},\beta)P(\{\alpha_k\}). \tag{A.1}$$

Here $P(n) = \sum_k \mathcal{L}(n|\{\alpha_k\},\beta,m)P(\{\alpha_k\})$ is the total probability of detecting n photons and

$$\mathcal{L}(n|\alpha_k,\beta) = \mathrm{Tr}\Big[|\alpha_k\rangle\langle\alpha_k|\hat{D}^\dagger(\beta)|n\rangle\langle n|\hat{D}(\beta)\Big] \tag{A.2}$$

$$= |\langle n|\hat{D}(\beta)|\alpha_k\rangle|^2 \tag{A.3}$$

$$= \frac{(|\alpha_k-\beta|^2)^n}{n!}e^{-|\alpha_k-\beta|^2}, \tag{A.4}$$

is the likelihood of detecting n photons from the state $|\alpha_k\rangle$ given the displacement amplitude β. The posterior distribution for step j then becomes the prior probability distribution for next adaptive step $j+1$, and this recursive procedure repeats until the entire state has been measured. Thus, the *final* posterior probability of each state given the displacements $\{\beta_j\}_L$ and detection history $\{n_j\}_L$ is:

$$P(\{\alpha_k\}|\{n_j\}_L,\{\beta_j\}_L) = \mathcal{N}P(\{\alpha_k\})\prod_{j=1}^{L}\mathcal{L}(n_j|\{\alpha_k\},\beta_j), \tag{A.5}$$

where $\mathcal{N}$ is a normalization factor and $P(\{\alpha_k\})$ is the initial prior probability distribution. This final posterior distribution is then used to infer the true input state depending on the task the receiver is performing. The hypothesis for the actual input in state discrimination is the state with the largest posterior probability. Alternatively, the estimate for the input state in phase estimation is obtained using Eq. (5.7) in Ch. 5 by integration over $P(\{\alpha_k\}|\{n_j\}_L,\{\beta_j\}_L)$, similar to a Bayesian estimator. Furthermore, the measurement outcome may be inconclusive, where the receiver guesses *none* of the possible input states as in Section 2.3.

A.2 Detailed Experimental Setup

Figure A.1 shows the detailed experimental setup used for most of the experiments described in the main thesis. Optical pulses of light are generated by a continuous-wave

Helium-Neon (HeNe) laser at 633 nm and a pulsed acousto-optic modulator (AOM) by selecting only the first order diffraction mode. The pulses then enter a Mach-Zender interferometer where a 50/50 beam splitter (BS) splits the pulse into the signal and local oscillator (LO) arms of the interferometer. We use a half-wave plate (HWP) and quarter-wave plate (QWP) to control the polarization of the light in each arm. A variable attenuator (Att.) and a fiber coupled phase modulator (PM) prepare the input state (signal) with fixed mean photon number and phase. A second PM and a fiber coupled amplitude modulator (AM) prepare the LO field used to implement the displacement operation. The signal field and LO field then interfere on a 99/1 fiber BS (purple box), which is where the displacement operation takes place. The light exits the interferometer and is detected either by a biased detector (Det.) during calibration using a flip-mirror, or by an avalanche photo-diode (APD) during the experiment to detect single photons.

The experiments are usually operated at a $\approx 50\%$ duty cycle and we maintain a constant relative phase between the signal and LO by actively locking the relative

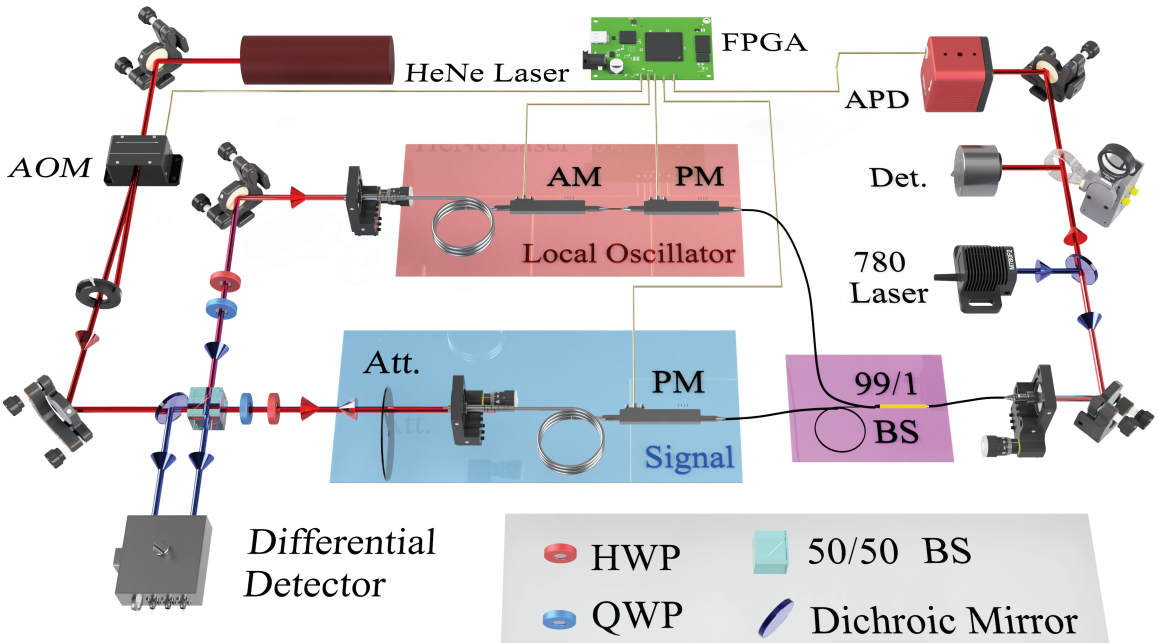

Figure A.1: Experimental setup, see main text for details.

phase in the interferometer between each experiment. We use a 780 nm laser to lock the interferometer phase, and frequency stabilize the 780 nm laser to rubidium using the saturated absorption spectroscopy locking technique. The 780 nm light propagates backwards through the interferometer and is detected by a differential detector (DD). We use dichroic mirrors which have $\approx 99\%$ transmission at 633 nm with a ≈ 1 nm bandwidth to separate the 633 nm light and 780 nm light. The output of the DD is fed into a PID controller which has an output that is connected to a piezo-electric crystal (PZT) in the signal arm. This feedback loop actively maintains a constant relative phase in the interferometer and holds its current value during each experimental pulse. The quantum efficiency of our APD (Laser Components COUNT-T250) is $\eta_{APD} \approx 0.82$ and we achieve an overall efficiency of $\eta \approx 0.72$ depending on the experiment. A "blanking" period of least ≈ 160 ns is required for adaptive strategies where the output of the APD is temporarily discarded. The blanking period corresponds to the time in between when the FPGA begins to apply new voltage values to the modulators and when the light propagates from the modulators to the APD. This blanking period is required for each adaptive step and artificially reduces the overall detection efficiency depending on how much blanking is implemented.

The experiments are timed and controlled by a field-programmable gate array (FPGA), which also registers all photon detections and passes data to a computer for processing. We use an Altera Cyclone IV (EP4CE55F23C8N) from Opal Kelly (ZEM4310) which has a 50 MHz base clock, 160 output pins, and 128 MiB external memory. Experiments are run at different repetition rates depending on the complexity of the measurement strategy and the amount of data saved for a single experiment. Each modulator is controlled by an 8-bit digital-to-analog converter (DAC) which allows for the phase and magnitude of each field (signal and LO) to take on 256 possible values. We also use a series of multiplexers as high bandwidth switches to turn the modulation on and off, or switch between two different modulation voltages.

A.3 Experimental Limitations

One critical experimental parameter is the interference visibility ξ of the displacement operation, which characterizes how well the input state and local oscillator field can interfere. The two fields will destructively interfere when they have the same magnitude and a π phase difference, resulting in a minimum intensity I_{min}, which is ideally equal to zero. A phase difference of zero will result in the fields constructively interfering to give a maximum intensity I_{max}. The interference visibility can then be obtained as $\xi = (I_{max} - I_{min})/(I_{max} + I_{min})$ where the background level has already been subtracted from both I_{max} and I_{min}. Our experiment achieves a visibility of $\xi \approx 0.998$, corresponding to $I_{min}/I_{max} \approx 1/1000$. Another important experimental parameter is the dark count rate ν of the photon detector, which corresponds to "detecting" a photon which did not come from the input light. This rate includes the dark counts intrinsic to the detector itself as well as detections from background light and can be modeled as a Poisson process with rate ν, and we achieve $\nu \approx 160/\text{s}$.

A realistic detector will not have the ability to resolve infinitely many photons, thus it is important to consider having a finite photon number resolution (PNR). The resolution of a detector is denoted as PNR(m), where m describes how many photons a detector can individually resolve before become a threshold detector. For example, the outcomes from a PNR(3) detector are $\{0, 1, 2, 3+\}$ where 3+ represents detecting three or more photons. Instead of the 3+ outcome being represented by $|3\rangle\langle 3|$, it is given by $\hat{I} - |0\rangle\langle 0| - |1\rangle\langle 1| - |2\rangle\langle 2|$ such that $\sum_n \hat{\Pi}_n(\beta) = \hat{I}$. As discussed throughout this thesis, the ability to count photons is a critical aspect in surpassing the limits of conventional Gaussian measurements. Furthermore, PNR increases the robustness of non-Gaussian receivers to noise and imperfections in the measurement as well as the communication channel.

Appendix B

Experimental Details

B.1 Implementing Photon Number Resolution

We allow the receiver to have photon number resolution (PNR) greater than PNR(1) in order to demonstrate many of the experiments described in the main thesis. However, we use an avalanche photo-diode (APD) in the experiment, which can only detect a single photon at any moment in time. Thus, PNR refers to the ability to detect multiple individual photons within a single adaptive step. This is in contrast to true PNR detectors, such as transition-edge sensors [136], and our *effective* PNR relies on two things. First, the temporal statistics of coherent states are such that photons are randomly distributed in time, as opposed to bunched (e.g. thermal) or anti-bunched (e.g. squeezed) light [8]. Second, we set our experimental repetition rate such that the time duration of each adaptive step T/L is much greater than the dead-time T_{dead} of the APD ($T_{dead} \approx 50$ ns). The dead-time refers to the time immediately after a detection event where the APD "resets" and cannot detect another photon. As long as the ratio of the dead-time to the adaptive step duration LT_{dead}/T is small, then the probability that there is a photon present during the dead-time is negligible. Thus, PNR for our experiment refers to the total number of individual photons detected over the course of a single adaptive step, each detected one at a time.

B.2 APD After-Pulsing Model

After-pulsing is an undesired effect where an avalanche photo-diode (APD) outputs more than one "click" when just a single photon is present [263]. When the APD detects a single (real) photon, it triggers an electron avalanche and results in the APD outputting a macroscopic electrical pulse followed by a "dead-time" where no photons may be detected. However, electrons can become trapped after a detection event and the trapped electrons are released upon recharging of the APD after the dead-time and trigger a second "fake" detection, or after-pulse, with probability P_{AP}. Furthermore, this effect can again trigger itself if another electron is trapped during the avalanche due to the "fake" detection, causing three "clicks" from a single real photon, and so on. This effect of multiple output pulses following the detection of a single real photon is known as after-pulsing. Although complex after-pulsing models have been developed [264], we model this process in a simple and intuitive way. This simple model works well for our "effective" PNR detector where the precise detection time of a photon is not recorded and we use just the information of whether or not photons were detected within a pulse for non-adaptive strategies, or within a single adaptive step. If the measurement strategy receives a single "click", then that detection must be from a single real photon. If the there are two "clicks" then that could have resulted from either two real photons, or a single real photon and a single after-pulse. Thus, the probability of receiving two "clicks" is modified as: $\tilde{P}(2) = P(2) + P(1)P_{AP}$ where $P(n)$ is the probability of there being n real photons in the state $|\alpha_k - \beta\rangle$ the APD is measuring. In a similar way, there are four ways to receive three "clicks" comprising of real (R) photons and after-pulses (A): $\{RRR, RRA, RAR, RAA\}$. Thus, the probability of receiving three "clicks" is $\tilde{P}(3) = P(3) + 2P(2)P_{AP} + P(1)P_{AP}^2$. Note it is impossible to have an after-pulse without first detecting a real photon, i.e. the outcomes $\{ARR, AAR, ARA, AAA\}$ are impossible, and it is also impossible to have an after-pulse given a single "click". This procedure can be extended to higher order photon detection events but when incorporating after-pulsing into our experimental model, we only use $\tilde{P}(2)$ and $\tilde{P}(3)$. We determined the after-pulsing probability for our

detector to be $P_{AP} = 1.10 \times 10^{-2}$ using time-delayed photon counting measurements.

B.3 Limitations of the Optimally Inconclusive Measurement

We implement the optimally inconclusive measurement in Chapter 2.3 using the optimal waveform $u(t)$ for the displacement operation. The ratio of the LO mean photon number to the input mean photon number $R = |u(t)|^2/|\alpha|^2$ is ideally infinite [53,62,90]. However, the maximum value of R that we experimentally implement is $R = 50$ due to the finite extinction ratio of ≈20 dB of the amplitude modulator (AM) in the LO arm of the interferometer, which corresponds to a non-ideal AM fringe visibility of ≈ 0.980. We calibrate the LO intensity such that the maximum transmission through the LO AM corresponds to $|\beta|^2 = \max(|u(t)|^2)$ and then a particular voltage is applied such that $|\beta|^2 = |\alpha|^2$ for calibration. However, the properties of the AM slightly drift over time and the AM changes the polarization of the light (thus degrading the interference visibility) when very close to the bottom of the AM fringe, i.e. at large modulation values due to large R. Therefore, we fix $R = 50$ to reliably implement the optimally inconclusive strategies even though an extinction ratio of 20 dB implies that $R = 100$ is possible.

Figure B.1(a) shows the effect of different values of R on the performance of the optimally inconclusive receiver with our experimental parameters. The blue, orange, and yellow lines show the error probability P_E as a function of the inconclusive probability P_I for strategies with $|\alpha|^2 = 0.2, 0.4$, and 0.6, respectively. The dashed lines shows the results for $R = \infty$ and the solid lines show the results for $R = 50$ as in our implementation. We find that the main effect of a finite value of R is that the receiver is unable to implement the optimally inconclusive measurement near the "middle" of each curve in Fig. B.1(a), e.g. $0.31 < P_I < 0.45$ for $|\alpha|^2 = 0.2$. We characterize this effect as a gap g in the $\{P_I, P_E\}$ curves as $g = (\Delta P_I)^2 + (\Delta P_E)^2$

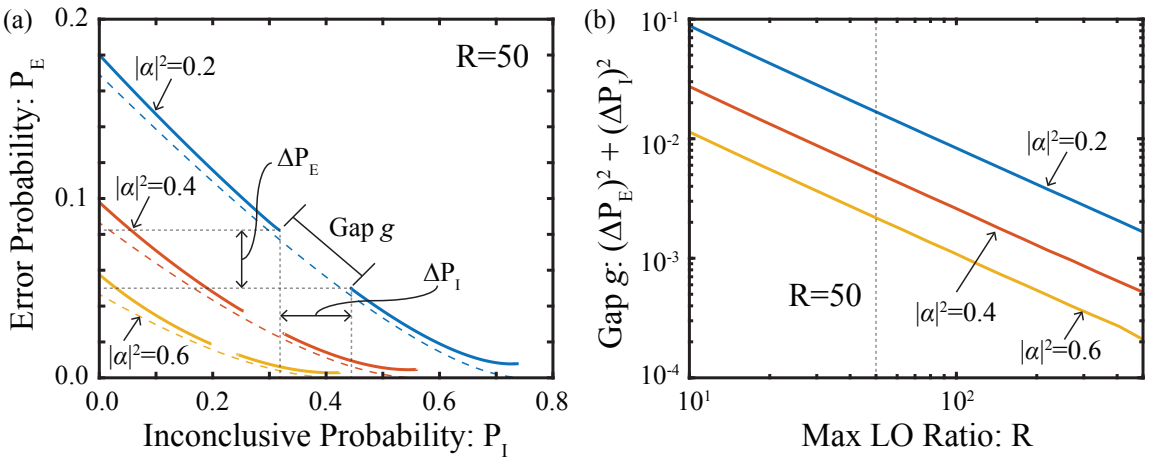

Figure B.1: (a) Expected results for the optimal inconclusive measurement with imperfections and a fixed maximum LO ratio $R = 50$ (solid) compared to an ideal measurement (except $\eta = 0.72$) with $R = \infty$ (dashed). The finite value of R causes a gap g in the curves near the middle. (b) Size of the gap g as a function of R for the different mean photon numbers. Fitting reveals that this gap scales as: $g \propto 1/R$ for all $|\alpha|^2$.

where ΔP_I and ΔP_E are shown by the dashed gray lines. Figure B.1(b) shows how g scales with R for $|\alpha|^2$=0.2, 0.4, and 0.6 in blue, orange, and yellow, respectively. We find that the gap scales as $g \propto 1/R$ for all the mean photon numbers from curve fitting.

B.4 Phase Estimation Algorithm

Part of the phase estimation algorithm discussed in Section 5.3 reconstructs the likelihood function $\mathcal{L}(d_j|\phi, |\beta|_{opt}, \theta_{opt})$ given a particular detection d_j, displacement magnitude $|\beta|_{opt}$, and displacement phase θ_{opt}. The algorithm leverages both the symmetries and circular properties of the likelihood functions for photon counting to efficiently reconstruct $\mathcal{L}(d_j|\phi, |\beta|_{opt}, \theta_{opt})$. The FPGA performs a "mirror & select" operation in order to effectively implement a circular shift of $\mathcal{L}(d_j|\phi, |\beta|_{opt}, 0)$ by the phase of the displacement θ_{opt} to obtain $\mathcal{L}(d_j|\phi, |\beta|_{opt}, \theta_{opt})$, which is easy to implement in MATLAB but more difficult in an FPGA. First, the algorithm retrieves half of the likelihood function $\mathcal{L}(d_j|\phi, |\beta|_{opt}, 0)$ from the internal FPGA memory corresponding

to the range $\phi = [0, \pi]$ given d_j and $|\beta|_{opt}$, but with no displacement phase. Next, this half of a single likelihood function is mirrored to obtain $\mathcal{L}(d_j|\phi, |\beta|_{opt}, 0)$ over the phase range $\phi = [0, 2\pi]$, and then mirrored again to obtain the likelihood over $\phi = [0, 4\pi]$, which is simply two copies of $\mathcal{L}(d_j|\phi, |\beta|_{opt}, 0)$ for $\phi = [0, 2\pi]$ side by side. Finally, the FPGA selects the phase range $\phi = [2\pi - \theta_{opt}, 4\pi - \theta_{opt}]$ of the doubly-mirrored likelihood function. Thus effectively implementing a circular shift such that the selected region corresponds to the phase range $\phi = [0, 2\pi]$ of the complete likelihood function $\mathcal{L}(d_j|\phi, |\beta|_{opt}, \theta_{opt})$ for the detection result d_j given $\beta_{opt} = |\beta_{opt}|e^{i\theta_{opt}}$.

The phase estimation algorithm calculates the posterior distribution $P^{(j)}_{post}(\phi)$ and also implements a quasi-normalization of $P^{(j)}_{post}(\phi)$ in order to maintain enough precision for each probability in the distribution. This normalization step multiplies the entire distribution by an amount equal to 2^p, where p is an integer which corresponds to the number of leading zeros in the binary representation of $P^{(j)}_{post}(\hat{\phi}_{MAP})$. We multiply $P^{(j)}_{post}(\phi)$ by a power of two because this is equivalent to shifting the binary representation of the probabilities p bits to the left. Shifting the posterior probabilities is an easy operation to efficiently implement in the FPGA and allows for maintaining the maximum amount of precision in the probability distributions. For example, if originally $P^{(j)}_{post}(\hat{\phi}_{MAP}) \approx 0.174$ then it has a binary representation of $P^{(j)}_{post}(\hat{\phi}_{MAP}) = [0010110010]$, and there are two leading zeros ($p = 2$) so the entire posterior distribution is multiplied by $2^p = 2^2$ such that now $P^{(j)}_{post}(\hat{\phi}_{MAP}) = [1011001000] = 0.695$. The quasi-normalization forces the maximum posterior probability $P^{(j)}_{post}(\hat{\phi}_{MAP})$ to be in between 0.5 and 1.0 and is completed in just a few nanoseconds. This quasi-normalization also allows the FPGA to maintain an accurate representation of the true distribution, since probabilities with values less than $1/2^{10} \approx 0.001$ are set equal to zero due to only having 10 bits of precision.

The algorithm also requires the prior distribution $P^{(j)}(\phi)$ to be *actually* normalized to calculate the variance σ^2 and skewness μ_3. The FPGA efficiently calculates these moments by first constructing a coarse-grained version $\tilde{P}^{(j)}(\phi)$ of the quasi-normalized prior distribution $P^{(j)}(\phi)$, which is centered at $\hat{\phi}_{MAP}$ and consists of 31 probabilities.

The algorithm obtains the values of $\tilde{P}^{(j)}(\phi)$ by starting from $\hat{\phi}_{MAP}$ and taking the values of $P^{(j)}(\phi)$ in steps of 8 moving outwards from $\hat{\phi}_{MAP}$. Thus, the 31 probabilities in $\tilde{P}^{(j)}(\phi)$ consist of $\hat{\phi}_{MAP}$ and 15 points on either side of $\hat{\phi}_{MAP}$. For example, if the index of $\hat{\phi}_{MAP}$ is $i = 100$ ($\hat{\phi}_{MAP} = \phi_{100}$) then then phase values used for $\tilde{P}^{(j)}(\phi)$ are:

$$\{\phi_{-20}, \phi_{-12}, \phi_{-4}, \phi_4, ..., \phi_{92}, \phi_{100}, \phi_{108}, ..., \phi_{220}\}, \tag{B.1}$$

such that

$$\tilde{P}^{(j)}(\phi) = \{P^{(j)}(\phi_{-20}), ..., P^{(j)}(\phi_{92}), P^{(j)}(\phi_{100}), P^{(j)}(\phi_{108}), ..., P^{(j)}(\phi_{220})\}. \tag{B.2}$$

Due to the circular properties of the phase distributions, i.e. $P^{(j)}(\phi_i) = P^{(j)}(\phi_{i+256})$ since $P^{(j)}(\phi) = P^{(j)}(\phi + 2\pi)$, the phase values in Eq. (B.1) are equivalent to:

$$\{\phi_{236}, \phi_{244}, \phi_{252}, \phi_4, ..., \phi_{92}, \phi_{100}, \phi_{108}, ..., \phi_{220}\}. \tag{B.3}$$

Thus, the coarse-grained prior probability distribution $\tilde{P}^{(j)}(\phi)$ is obtained as:

$$\tilde{P}^{(j)}(\phi) = \{P^{(j)}(\phi_{236}), ..., P^{(j)}(\phi_{92}), P^{(j)}(\phi_{100}), P^{(j)}(\phi_{108}), ..., P^{(j)}(\phi_{220})\}. \tag{B.4}$$

This coarse grained prior distribution $\tilde{P}^{(j)}(\phi)$ is then used to efficiently calculate the variance σ^2 and skewness μ_3. The calculation is done by first obtaining the normalization factor $\mathcal{N} = \sum_k \tilde{P}^{(j)}(\phi_k)$ and then instead of immediately normalizing $\tilde{P}^{(j)}(\phi)$, we calculate the variance and skewness as:

$$\sigma^2 = \frac{1}{\mathcal{N}} \times \left(\sum_{k=1}^{31} \phi_k^2 \tilde{P}^{(j)}(\phi_k) \right), \tag{B.5}$$

$$\mu_3 = \frac{1}{\mathcal{N}} \times \left(\sum_{k=1}^{31} \phi_k^3 \tilde{P}^{(j)}(\phi_k) \right), \tag{B.6}$$

where the summation terms are calculated prior to division by $\mathcal{N}$ in the FPGA. This allows the algorithm to efficiently calculate σ^2 and μ_3 using only a single division operation (which takes ≈ 200 ns) as opposed to many divisions to obtain a normalized distribution first. We note that the quasi-normalization term 2^p will actually be canceled out due to the fact that it appears in both $\tilde{P}^{(j)}(\phi)$ and $\mathcal{N}$. Thus, the value of p does not actually influence σ^2 and μ_3, and this procedure can be applied every adaptive measurement step.

Appendix C

Genetic Optimization Algorithm

We implement an evolutionary optimization strategy in Ch. 3.2 to optimize the alternative receiver based on the LSTM/NN architecture. Evolutionary optimization algorithms rely on maintaining a pool of candidate solutions $\{V_i\}$ where V_i is an n-dimensional vector. We aim to find solutions which are vectors of real numbers corresponding to all the weights and biases of the LSTM/NN architecture. Thus, we use an evolutionary algorithm referred to as differential evolution (DE), which can be used for functions which are discontinuous, noisy, or even time-dependent [149,150]. A single iteration of the DE algorithm follows two basic steps: evaluation and evolution. The first step is evaluating all the candidate solutions to obtain their respective *fitness* F_i, which in our case is the error probability $P_E^{(i)}$. The value of $P_E^{(i)}$ is sampled through Monte Carlo simulations using the LSTM/NN receiver to implement state discrimination. The second step is evolution, where the algorithm constructs a new pool of solutions from the current pool based on each candidates fitness, hence the term evolutionary optimization.

Algorithm 2 shows the DE algorithm we implement including how new candidate solutions are constructed. The left column is the main DE algorithm and the right column shows the evolution (EVOLVE$(\cdot)$) and evaluation (EVALUATE$(\cdot)$) functions used by the algorithm. Each candidate vector V_i is first initialized (left column, line

Algorithm 2 Differential Evolution (DE) Algorithm

```
1:  Begin DE algorithm
2:  Initial: {V_i} = {V_{i,initial}}, S_best = {}
3:  for i ← 1 to N do                                   ▷ Initial Fitness
4:      F_i ← EVALUATE(V_i, N_eval)
5:  end for
6:  F_best = min({F_i}), V_best = argmin({V_i})
7:
8:  loop
9:      S_best ← Best p percent of {V_i} based on {F_i}
10:     for i ← 1 to N do                               ▷ Evolve worst candidates
11:         V_i ← EVOLVE(S_best) if V_i ∉ S_best
12:         F_i ← EVALUATE(V_i, N_eval)
13:     end for
14:     if min(F) < F_best then                         ▷ Check possible best
15:         V_check ← candidate who has min({F_i})
16:         F' ← EVALUATE(V_check, N_check)
17:         if F' < F_best then
18:             F_best ← F'                             ▷ Update best fitness
19:             V_best ← V_check                        ▷ Update best candidate
20:         end if
21:     end if
22: end loop
```

```
1:  function EVOLVE(S_best)
2:      a, b, c, d ← Sample 4 candidates from S_best
3:      if rand(0,1)≤ p_mutate then
4:          y_new ← a + N_n(0, f_mutate)                ▷ Mutate
5:      else
6:          y_new ← a + f_cross(b − c)                  ▷ Crossover
7:      end if
8:      V_new ← d                                       ▷ Initialize new candidate
9:      if rand(0,1)≤0.5 then                           ▷ Evolve LSTM
10:         for j ← 1 to n_LSTM do
11:             if rand(0,1)< p_evolve then
12:                 V_new[j] ← y_new[j]
13:             end if
14:         end for
15:     else                                            ▷ Evolve NN
16:         for j ← n_LSTM + 1 to n do
17:             if rand(0,1)< p_evolve then
18:                 V_new[j] ← y_new[j]
19:             end if
20:         end for
21:     end if
22:     return V_new
23: end function
24:
25: function EVALUATE(V_i, N)
26:     Set LSTM/NN parameters to V_i
27:     P_E ← Evaluate with Monte Carlo with N
28:     return P_E
29: end function
```

2) and then the all candidates are evaluated (line 3 to 5). There are three steps for each iteration (line 9 to 21) in the differential evolution algorithm. First (line 9), the algorithm obtains the subset $S_{best} \subset \{V_i\}$ which represents the best p percent of all candidates based on F_i, where a smaller value of F_i is better. Each candidate not in S_{best} is then evolved (line 11) based on the candidates in S_{best}. Next, the fitness F_i each candidate V_i is evaluated (line 12) based on $N_{eval} = 10^3$ individual state discrimination measurements. Finally, if there is a candidate which has a fitness which is better than the current overall best value F_{best}, then it is checked/re-evaluated (line 15, 16) with $N_{check} = 10^4$ to reduce the effects of finite sampling. If the re-evaluated fitness F' is better than F_{best}, then V_{check} replaces the best known candidate V_{best} (line 17 to 20).

The algorithm evolves each candidate vector V_i not in S_{best} using the EVOLVE$(\cdot)$ function (right column, lines 1 to 23) by first randomly selecting four "good" candidates

$\{a, b, c, d\}$ from S_{best} (line 2). The type of evolution (mutation vs. crossover) is then randomly chosen with probability $p_{mutate} = 0.5$. Mutation (line 4) occurs by constructing a temporary candidate y_{new} and adding n-dimensional Gaussian noise with variance $f_{mutate} = 0.025$ to candidate a: $y_{new} = a + \mathcal{N}_n(0, f_{mutate})$. Crossover (line 6) occurs by constructing y_{new} as: $y_{new} = a + f_{cross}(b - c)$ where the strength is controlled by $f_{cross} = 0.25$. The next step is to initialize a new candidate vector $V_{new} = d$ (line 8) and then randomly choosing to update either the part of the vector corresponding to the LSTM parameters (line 9 to 14) or the part corresponding to the NN parameters (line 15 to 21) with equal probability. Each parameter $V_{new}[j]$ in the new candidate vector is replaced by the temporary candidates value $y_{new}[j]$ with probability $p_{evolve} = 0.5$ (line 12 or 18). This allows for randomly evolving a random subset of either the LSTM or NN parameters by either mutation or crossover. A candidate vector is evaluated using the EVALUATE$(\cdot)$ function (right column, lines 25 to 29). The first step (line 26) sets the LSTM/NN parameters to the candidate vector V_i to be evaluated. Then (line 27) the LSTM/NN is evaluated using Monte Carlo simulations of the state discrimination measurement with N individual experiments.

The final aspect of the optimization strategy is how each candidate vector is initialized. If each V_i randomly initialized, then the algorithm will in general fail to converge to a good solution even after thousands of iterations. We utilize another machine learning technique called transfer learning to address this issue of initialization. Transfer learning is when the network is trained on a related but different problem and then the result of that training is used as an initialization point for the original problem [151, 152]. We initially train the LSTM/NN using gradient descent to approximate a non-optimized phase estimation strategy (i.e. $M = 256$ states vs. $M = 4$ for QPSK) with only a few adaptive steps $L = 5$ to implement transfer learning. This technique allows the network to learn some form of adaptive Bayesian inference while remaining untrained for the problem we *actually* want to solve. We extract the "trained" parameters $\tilde{V}$ at the end of this initial training and initialize each candidate for the DE algorithm as: $V_i = \tilde{V} + \mathcal{N}_n(0, 0.1)$.

Bibliography

[1] Daniel Braun, Gerardo Adesso, Fabio Benatti, Roberto Floreanini, Ugo Marzolino, Morgan W. Mitchell, and Stefano Pirandola. Quantum-enhanced measurements without entanglement. *Rev. Mod. Phys.*, 90:035006, September 2018.

[2] C W Helstrom. *Quantum detection and estimation theory, Mathematics in Science and Engineering Vol. 123.* Academic Press, New York, 1976.

[3] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum-enhanced measurements: Beating the standard quantum limit. *Science*, 306(5700):1330–1336, 2004.

[4] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Advances in quantum metrology. *Nature Photon.*, 5:222–229, 2013.

[5] Christian Weedbrook, Stefano Pirandola, Raúl García-Patrón, Nicolas J. Cerf, Timothy C. Ralph, Jeffrey H. Shapiro, and Seth Lloyd. Gaussian quantum information. *Rev. Mod. Phys.*, 84:621–669, 2012.

[6] Alexander S. Holevo. *Probabilistic and Statistical Aspects of Quantum Theory.* Edizioni della Normale, 2011.

[7] Samuel L. Braunstein and Carlton M. Caves. Statistical distance and the geometry of quantum states. *Phys. Rev. Lett.*, 72:3439–3443, May 1994.

[8] Marlan O. Scully, Scully, and M. Suhail Zubairy. *Quantum Optics.* Cambridge University Press, September 1997.

[9] H Wiseman and G Milburn. *Quantum Measurement and Control.* Cambrige University Press, New York, 2010.

[10] Rafal Demkowicz-Dobrzański, Marcin Jarzyna, and Jan Kołodyński. Quantum limits in optical interferometry. *Progress in Optics*, 60:345–435, 2015.

[11] B. M. Escher, R. L. de Matos Filho, and L. Davidovich. General framework for estimating the ultimate precision limit in noisy quantum-enhanced metrology. *Nature Physics*, 7(5):406–411, May 2011.

[12] Carlton M. Caves. Quantum-mechanical noise in an interferometer. *Phys. Rev. D*, 23:1693–1708, Apr 1981.

[13] Carlton M. Caves and P. D. Drummond. Quantum limits on bosonic communication rates. *Rev. Mod. Phys.*, 66:481–537, 1994.

[14] J. A. Bergou, U. Herzog, and M. Hillery. Discrimination of quantum states. *Lect. Notes Phys.*, 649:417–465, 2004.

[15] Stephen M. Barnett and Sarah Croke. Quantum state discrimination. *Adv. Opt. Photon.*, 1, 2009.

[16] Anthony Chefles. Quantum state discrimination. *Contemporary Physics*, 41(6):401–424, 2000.

[17] Anthony Chefles and Stephen M. Barnett. Strategies for discriminating between non-orthogonal quantum states. *Journal of Modern Optics*, 45(6):1295–1302, 1998.

[18] V. Giovannetti, S. Guha, S. Lloyd, L. Maccone, J. H. Shapiro, and H. P. Yuen. Classical capacity of the lossy bosonic channel: The exact solution. *Phys. Rev. Lett.*, 92:027902, 2004.

[19] Peter van Loock, Norbert Lütkenhaus, W. J. Munro, and Kae Nemoto. Quantum repeaters using coherent-state communication. *Phys. Rev. A*, 78:062319, 2008.

[20] Saikat Guha. Structured optical receivers to attain superadditive capacity and the holevo limit. *Phys. Rev. Lett.*, 106:240502, 2011.

[21] Matteo Rosati, Andrea Mari, and Vittorio Giovannetti. Multiphase hadamard receivers for classical communication on lossy bosonic channels. *Phys. Rev. A*, 94:062325, 2016.

[22] A Klimek, M Jachura, W Wasilewski, and K Banaszek. Quantum memory receiver for superadditive communication using binary coherent states. *Journal of Modern Optics*, 63(20):2074–2080, 2016.

[23] Konrad Banaszek, Ludwig Kunz, Michal Jachura, and Marcin Jarzyna. Quantum Limits in Optical Communications. *Journal of Lightwave Technology*, 38(10):2741 – 2754, 2020.

[24] C H Bennet and . Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing*, Malvern Physics Series, page 175, Bangalore, 1984.

[25] C H Bennett. Quantum cryptography using any two nonorthogonal states. 68:3121–3124, 1992.

[26] B. Huttner, N. Imoto, N. Gisin, and T. Mor. Quantum cryptography with coherent states. *Phys. Rev. A*, 51:1863–1869, 1995.

[27] Frédéric Grosshans and Philippe Grangier. Continuous variable quantum cryptography using coherent states. *Phys. Rev. Lett.*, 88:057902, 2002.

[28] Ch. Silberhorn, T. C. Ralph, N. Lütkenhaus, and G. Leuchs. Continuous variable quantum cryptography: Beating the 3 db loss limit. *Phys. Rev. Lett.*, 89:167901, 2002.

[29] F. Grosshans, G. van Assche, J. Wenger, R. Tualle-Brouri, N. J. Cerf, and P. Grangier. Quantum key distribution using gaussian-modulated coherent states. *Nature*, 421:238, 2003.

[30] N Gisin, G Ribordy, W Tittel, and H Zbinden. Quantum cryptography. *Rev. Mod. Phys.*, 74:145–195, 2002.

[31] D. Sych and G. Leuchs. Coherent state quantum key distribution with multi letter phase-shift keying. *New J. of Phys.*, 12(5):053019, 2010.

[32] Rui Han, Gerd Leuchs, and Markus Grassl. Residual and destroyed accessible information after measurements. *arXiv.org*, quant-ph/arXiv:1710.11586, 2017.

[33] Rui Han, János A. Bergou, and Gerd Leuchs. Near optimal discrimination of binary coherent signals via atom–light interaction. *New Journal of Physics*, 20(4):043005, April 2018.

[34] W J Munro, K Nemoto, and T P Spiller. Weak nonlinearities: a new route to optical quantum computation. *New J. of Phys.*, 7(1):137, 2005.

[35] P van Loock. Optical hybrid approaches to quantum information. *Laser & Photon. Rev.*, 5:167–200, 2011.

[36] Kae Nemoto and W. J. Munro. Nearly deterministic linear optical controlled-not gate. *Phys. Rev. Lett.*, 93:250502, 2004.

[37] T. C. Ralph, A. Gilchrist, G. J. Milburn, W. J. Munro, and S. Glancy. Quantum computation with optical coherent states. *Phys. Rev. A*, 68:042319, 2003.

[38] U Leonhardt. *Measuring the Quantum State of Light*. Cambridge University Press, Cambridge, 1997.

[39] Valeria Cimini, Marta Mellini, Giordano Rampioni, Marco Sbroscia, Livia Leoni, Marco Barbieri, Marco Barbieri, Ilaria Gianani, and Ilaria Gianani. Adaptive tracking of enzymatic reactions with quantum light. *Optics Express*, 27(24):35245–35256, November 2019.

[40] Andrea Crespi, Mirko Lobino, Jonathan C. F. Matthews, Alberto Politi, Chris R. Neal, Roberta Ramponi, Roberto Osellame, and Jeremy L. O'Brien. Measuring protein concentration with entangled photons. *Applied Physics Letters*, 100(23):233704, June 2012.

[41] Michael A. Taylor, Jiri Janousek, Vincent Daria, Joachim Knittel, Boris Hage, Hans-A. Bachor, and Warwick P. Bowen. Biological measurement beyond the quantum limit. *Nature Photonics*, 7(3):229–233, March 2013.

[42] Takafumi Ono, Ryo Okamoto, and Shigeki Takeuchi. An entanglement-enhanced microscope. *Nature Communications*, 4(1):2426, September 2013.

[43] H. Noguchi, H. Kubo, T. Mori, T. Sato, and H. Sanada. Signal phase estimation for measurement of respiration waveform using a microwave Doppler sensor. In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 6740–6743, July 2013. ISSN: 1558-4615.

[44] S. Danilin, A. V. Lebedev, A. Vepsäläinen, G. B. Lesovik, G. Blatter, and G. S. Paraoanu. Quantum-enhanced magnetometry by phase estimation algorithms with a single artificial atom. *npj Quantum Information*, 4(1):1–8, June 2018.

[45] N. M. Nusran and M. V. Gurudev Dutt. Optimizing phase-estimation algorithms for diamond spin magnetometry. *Physical Review B*, 90(2):024422, July 2014.

[46] J. R. Maze, P. L. Stanwix, J. S. Hodges, S. Hong, J. M. Taylor, P. Cappellaro, L. Jiang, M. V. Gurudev Dutt, E. Togan, A. S. Zibrov, A. Yacoby, R. L. Walsworth, and M. D. Lukin. Nanoscale magnetic sensing with an individual electronic spin in diamond. *Nature*, 455(7213):644–647, October 2008.

[47] [. . .]J. Zweizig J. Aasi, J. Abadie. Enhanced sensitivity of the ligo gravitational wave detector by using squeezed states of light. *Nat. Photonics*, 7:613, 2013.

[48] Matteo G. A. Paris. Quantum estimation for quantum technology. *International Journal of Quantum Information*, 07(supp01):125–137, January 2009.

[49] D. Dieks. Overlap and distinguishability of quantum states. *Phys. Lett. A*, 126:303–306, 1988.

[50] K. Kikuchi and S. Tsukamoto. Evaluation of sensitivity of the digital coherent receiver. *Lightwave Technology, Journal of*, 26:1817–1822, 2008.

[51] F. E. Becerra, J. Fan, G. Baumgartner, J. Goldhar, J. T. Kosloski, and A. Migdall. Experimental demonstration of a receiver beating the standard quantum limit for multiple nonorthogonal state discrimination. *Nature Photonics*, 7:147–152, 2013.

[52] Kenji Tsujino, Daiji Fukuda, Go Fujii, Shuichiro Inoue, Mikio Fujiwara, Masahiro Takeoka, and Masahide Sasaki. Quantum receiver beyond the standard quantum limit of coherent optical communication. *Phys. Rev. Lett.*, 106:250503, 2011.

[53] Robert L. Cook, Paul J. Martin, and J. M. Geremia. Optical coherent state discrimination using a closed-loop quantum measurement. *Nature*, 446:774–777, 2007.

[54] C R Müller, M A Usuga, C Wittmann, M Takeoka, Ch Marquardt, U L Andersen, and G Leuchs. Quadrature phase shift keying coherent state discrimination via a hybrid receiver. *New J. of Phys.*, 14(8):083009, 2012.

[55] Shuro Izumi, Jonas S. Neergaard-Nielsen, Shigehito Miki, Hirotaka Terai, and Ulrik L. Andersen. Experimental Demonstration of a Quantum Receiver Beating the Standard Quantum Limit at Telecom Wavelength. *Physical Review Applied*, 13(5):054015, 2020.

[56] R J Glauber. The quantum theory of optical coherence. 130:2529, 1963.

[57] M A Nielsen and I L Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge England, 2000.

[58] K. Kikuchi. Fundamentals of coherent optical fiber communications. *Lightwave Technology, Journal of*, 34:157–179, 2016.

[59] K Kato, M Osaki, M Sasaki, and O Hirota. Quantum detection and mutual infromation for qam and psk signals. *IEEE Trans. On Comm.*, 47(2):248, 1999.

[60] Kentaro Wakui Mikio Fujiwara Kazuhiro Ema Shuro Izumi, Masahiro Takeoka and Masahide Sasaki. Projective measurement onto arbitrary superposition of weak coherent state bases. *Scientific Reports*, (8):2999, 2018.

[61] Rui Han, Gerd Leuchs, and János A. Bergou. Helstrom measurement: A nondestructive implementation. *Physical Review A*, 101(3):032103, 2020.

[62] S J Dolinar. An optimum receiver for the binary coherent state quantum channel. Research Laboratory of Electronics, MIT, Quarterly Progress Report No. 111 (1973), p. 115.

[63] Kenji Nakahira and Tsuyoshi Sasaki Usuda. Optimal receiver for discrimination of two coherent states with inconclusive results. *Physical Review A*, 86(5):052323, 2012.

[64] M.T. DiMario and F.E. Becerra. Experimental demonstration of optimally inconclusive measurements for binary coherent states. *In preparation*, 2021.

[65] Y. C. Eldar and G. D. Forney. On quantum detection and the square-root measurement. *IEEE Transactions on Information Theory*, 47(3):858–872, 2001.

[66] Masashi Ban, Keiko Kurokawa, Rei Momose, and Osamu Hirota. Optimum measurements for discrimination among symmetric quantum states and parameter estimation. *International Journal of Theoretical Physics*, 36(6):1269–1288, 1997.

[67] Nicola Dalla Pozza and Gianfranco Pierobon. Optimality of square-root measurements in quantum state discrimination. *Physical Review A*, 91(4):042334, 2015.

[68] Ezra Ip, Alan Pak Tao Lau, Daniel J. F. Barros, and Joseph M. Kahn. Coherent detection in optical fiber systems. *Opt. Express*, 16(2):753–791, 2008.

[69] Masahiro Takeoka and Masahide Sasaki. Discrimination of the binary coherent signal: Gaussian-operation limit and simple non-gaussian near-optimal receivers. *Phys. Rev. A*, 78:022320, 2008.

[70] Ranjith Nair, Brent J. Yen, Saikat Guha, Jeffrey H. Shapiro, and Stefano Pirandola. Symmetric m-ary phase discrimination using quantum-optical probe states. *Phys. Rev. A*, 86:022306, 2012.

[71] N. J. Cerf V. Giovannetti, R. García-Patrón and A. S. Holevo. Ultimate classical communication rates of quantum optical channels. *Nature Photonics*, 8:796–800, 2014.

[72] A. Klimek K. Banaszek M. G. A. Paris M. Jarzyna, V. Lipińska. Phase noise in collective binary phase shift keying with hadamard words. *Opt. Express*, 24(2):1693–1698, 2016.

[73] Jacopo Trapani, Berihu Teklu, Stefano Olivares, and Matteo G. A. Paris. Quantum phase communication channels in the presence of static and dynamical phase diffusion. *Phys. Rev. A*, 92:012317, July 2015.

[74] Alan Pak Tao Lau and Joseph M. Kahn. Signal Design and Detection in Presence of Nonlinear Phase Noise. *Journal of Lightwave Technology*, 25(10):3008–3016, October 2007.

[75] K. Banaszek M. Jarzyna and R. Demkowicz-Dobrzański. Dephasing in coherent communication with weak signal states. *J. Phys. A: Math. Theor.*, 47:275302, 2014.

[76] J. Shapiro. Quantum noise and excess noise in optical homodyne and heterodyne receivers. *IEEE Journal of Quantum Electronics*, 21(3):237–250, March 1985.

[77] J. G. Proakis. *Digital Communications, 4th Ed.* McGraw-Hill, New York, 2000.

[78] H. M. Wiseman and R. B. Killip. Adaptive single-shot phase measurements: The full quantum theory. *Phys. Rev. A*, 57:2169–2185, March 1998.

[79] Matteo G.A. Paris. Displacement operator by beam splitter. *Physics Letters A*, 217:78–80, 1996.

[80] M. T. DiMario and F. E. Becerra. Robust measurement for the discrimination of binary coherent states. *Phys. Rev. Lett.*, 121(2):023603, 2018.

[81] M. T. DiMario, L. Kunz, K. Banaszek, and F. E. Becerra. Optimized communication strategies with binary coherent states over phase noise channels. *npj Quantum Information*, 5:65, 2019.

[82] M. T. DiMario, E. Carrasco, R. A. Jackson, and F. E. Becerra. Implementation of a single-shot receiver for quaternary phase-shift keyed coherent states. *J. Opt. Soc. Am. B*, 35:568–574, 2018.

[83] A. R. Ferdinand, M. T. DiMario, and F. E. Becerra. Multi-state discrimination below the quantum noise limit at the single-photon level. *npj Quantum Information*, 3:43, 2017.

[84] M.T. DiMario and F.E. Becerra. Finding optimized non-gaussian measurements with black-box reinforcement learning. *In preparation*, 2021.

[85] M. T. DiMario and F. E. Becerra. Phase tracking for sub-shot-noise-limited receivers. *Phys. Rev. Research*, 2:023384, 2020.

[86] M. T. DiMario and F. E. Becerra. Channel-noise tracking for sub-shot-noise-limited receivers with neural networks. *Physical Review Research*, 3(1):013200, March 2021.

[87] M.T. DiMario and F.E. Becerra. Single-Shot Non-Gaussian Measurements for Optical Phase Estimation. *Physical Review Letters*, 125(12):120505, Sept. 2020.

[88] M Rodriguez-Garcia, M T DiMario, P Barberis-Blostein, and F E Becerra. Adaptive single-shot phase estimation strategies for coherent states. *In preparation*, 2021.

[89] M.T. DiMario and F.E. Becerra. Maximizing mutual information with non-gaussian receivers and probabalistic shaping. *In preparation*, 2021.

[90] JM Geremia. Distinguishing between optical coherent states with imperfect detection. *Physical Review A*, 70(6):062303, 2004.

[91] R S Kennedy. A near-optimum receiver for the binary coherent state quantum channel. Research Laboratory of Electronics, MIT Technical Report No. 110 (1972), unpublished.

[92] Christoffer Wittmann, Masahiro Takeoka, Katiúscia N. Cassemiro, Masahide Sasaki, Gerd Leuchs, and Ulrik L. Andersen. Demonstration of near-optimal discrimination of optical coherent states. *Phys. Rev. Lett.*, 101:210501, 2008.

[93] G. Cariolaro and G. Pierobon. Performance of quantum data transmission systems in the presence of thermal noise. *IEEE Transactions on Communications*, 58(2):623–630, February 2010.

[94] R. Yuan, M. Zhao, S. Han, and J. Cheng. Kennedy Receiver Using Threshold Detection and Optimized Displacement Under Thermal Noise. *IEEE Communications Letters*, 24(6):1313–1317, June 2020.

[95] Ludwig Kunz, Matteo G. A. Paris, and Konrad Banaszek. Noisy propagation of coherent states in a lossy Kerr medium. *JOSA B*, 35(2):214–222, February 2018.

[96] K Banaszek. Optimal receiver for quantum cryptography with two coherent states. *Phys. Lett. A*, 253:12–15, 1999.

[97] S. J. van Enk. Unambiguous state discrimination of coherent states with linear optics: Application to quantum cryptography. *Phys. Rev. A*, 66:042313, 2002.

[98] Christoffer Wittmann, Ulrik L. Andersen, Masahiro Takeoka, Denis Sych, and Gerd Leuchs. Demonstration of coherent-state discrimination using a displacement-controlled photon-number-resolving detector. *Phys. Rev. Lett.*, 104:100505, 2010.

[99] Andrea Mari, Vittorio Giovannetti, and Alexander S. Holevo. Quantum state majorization at the output of bosonic gaussian channels. *Nature Commun.*, 5:3826, 2014.

[100] Matteo Bina, Alessia Allevi, Maria Bondani, and Stefano Olivares. Homodyne-like detection for coherent state-discrimination in the presence of phase noise. *Optics Express*, 25(9):10685–10692, May 2017.

[101] Giovanni Chesi, Stefano Olivares, and Matteo G. A. Paris. Squeezing-enhanced phase-shift-keyed binary communication in noisy channels. *Phys. Rev. A*, 97:032315, March 2018.

[102] Stefano Olivares, Simone Cialdi, Fabrizio Castelli, and Matteo G. A. Paris. Homodyne detection as a near-optimum receiver for phase-shift-keyed binary communication in the presence of phase diffusion. *Phys. Rev. A*, 87:050303, May 2013.

[103] Marco G. Genoni, Stefano Olivares, and Matteo G. A. Paris. Optical phase estimation in the presence of phase diffusion. *Phys. Rev. Lett.*, 106:153603, April 2011.

[104] Marco G. Genoni, Stefano Olivares, Davide Brivio, Simone Cialdi, Daniele Cipriani, Alberto Santamato, Stefano Vezzoli, and Matteo G. A. Paris. Optical interferometry in the presence of large phase diffusion. *Phys. Rev. A*, 85:043817, April 2012.

[105] Marco G. Genoni Xian-Min Jin W. Steven Kolthammer M.S. Kim Animesh Datta Marco Barbieri Mihai D. Vidrighin, Gaia Donati and Ian A. Walmsley. Joint estimation of phase and phase diffusion for quantum metrology. *Nature Communications*, 3:1063, 2014.

[106] H. Yuen, R. Kennedy, and M. Lax. Optimum testing of multiple hypotheses in quantum detection theory. *IEEE Transactions on Information Theory*, 21(2):125–134, 1975.

[107] A. Peres. How to differentiate between non-otrhogonal states. *Phys. Lett. A*, 128:19, 1988.

[108] Gregg Jaeger and Abner Shimony. Optimal distinction between two non-orthogonal quantum states. *Physics Letters A*, 197(2):83–87, 1995.

[109] Jan Soubusta Lucie Bartůšková, Antonín Černoch and Miloslav Dušek. Programmable discriminator of coherent states: Experimental realization. *Phys. Rev. A*, 77:034306, 2008.

[110] Yonina C. Eldar. Mixed-quantum-state detection with inconclusive results. *Physical Review A*, 67(4):042309, 2003.

[111] H. Sugimoto, T. Hashimoto, M. Horibe, and A. Hayashi. Discrimination with error margin between two states: Case of general occurrence probabilities. *Physical Review A*, 80(5):052322, 2009.

[112] Jaromir Fiurasek and Miroslav Jezek. Optimal discrimination of mixed quantum states involving inconclusive results. *Physical Review A*, 67(1):012321, 2003.

[113] Antonio Assalini, Nicola Dalla Pozza, and Gianfranco Pierobon. Revisiting the Dolinar receiver through multiple-copy state discrimination theory. *Physical Review A*, 84(2):022342, 2011.

[114] A. Acín, E. Bagan, M. Baig, Ll. Masanes, and R. Muñoz-Tapia. Multiple-copy two-state discrimination with individual measurements. *Physical Review A*, 71(3):032338, 2005.

[115] Matthias Heid and Norbert Lütkenhaus. Efficiency of coherent-state quantum cryptography in the presence of loss: Influence of realistic error correction. *Physical Review A*, 73(5):052316, May 2006.

[116] Peter J. Winzer. High-spectral-efficiency optical modulation formats. *J. Lightwave Technol.*, 30, 2012.

[117] F. E. Becerra, J. Fan, and A. Migdall. Photon number resolution enables quantum receiver for realistic coherent optical communications. *Nat. Photonics*, 9, 2015.

[118] Jasminder S. Sidhu, Shuro Izumi, Jonas S. Neergaard-Nielsen, Cosmo Lupo, and Ulrik L. Andersen. Quantum receiver for quadrature phase-shift keying at the single photon level. *arXiv:2009.03339 [quant-ph]*, September 2020. arXiv: 2009.03339.

[119] Shuro Izumi, Masahiro Takeoka, Mikio Fujiwara, Nicola Dalla Pozza, Antonio Assalini, Kazuhiro Ema, and Masahide Sasaki. Displacement receiver for phase-shift-keyed coherent states. *Phys. Rev. A*, 86:042328, 2012.

[120] J T Kosloski. A kennedy receiver for optical quadrature phase shift keying. Ph.D., THE JOHNS HOPKINS UNIVERSITY, 2012, 257 pages; 3532695.

[121] Callum Croal, Christian Peuntinger, Bettina Heim, Imran Khan, Christoph Marquardt, Gerd Leuchs, Petros Wallden, Erika Andersson, and Natalia Korolkova. Free-space quantum signatures using heterodyne measurements. *Phys. Rev. Lett.*, 117:100503, 2016.

[122] Juan Miguel Arrazola, Markos Karasamanis, and Norbert Lütkenhaus. Practical quantum retrieval games. *Phys. Rev. A*, 93, 2016.

[123] J. M. Arrazola, P. Wallden, and E. Andersson. Multiparty quantum signature schemes. *Quantum Inf. Comput.*, 7, 2016.

[124] Roy S. Bondurant. Near-quantum optimum receivers for the phase-quadrature coherent-state channel. *Opt. Lett.*, 18(22):1896–1898, 1993.

[125] Shuro Izumi, Masahiro Takeoka, Kazuhiro Ema, and Masahide Sasaki. Quantum receivers with squeezing and photon-number-resolving detectors for m-ary coherent state discrimination. *Phys. Rev. A*, 87:042328, 2013.

[126] C R Müller and Ch Marquardt. A robust quantum receiver for phase shift keyed signals. *New Journal of Physics*, 17(3):032003, 2015.

[127] Ranjith Nair, Saikat Guha, and Si-Hui Tan. Realizable receivers for discriminating coherent and multicopy quantum states near the quantum limit. *Phys. Rev. A*, 89:032318, 2014.

[128] F. E. Becerra, J. Fan, G. Baumgartner, S. V. Polyakov, J. Goldhar, J. T. Kosloski, and A. Migdall. M-ary-state phase-shift-keying discrimination below the homodyne limit. *Phys. Rev. A*, 84:062324, 2011.

[129] C. Wei, X. Zhou, L. Wang, P. Tian, and L. Hanzo. Soft Iterative Quantum Receivers Approaching the Helstrom Limit Using Realistic Quantum Devices. *IEEE Access*, 6:10197–10207, 2018.

[130] Yuan Zuo, Ke Li, and Bing Zhu. 16-QAM Quantum Receiver with Hybrid Structure Outperforming the Standard Quantum Limit. *MATEC Web of Conferences*, 61:06008, 2016.

[131] Anthony Leverrier and Philippe Grangier. Unconditional security proof of long-distance continuous-variable quantum key distribution with discrete modulation. *Phys. Rev. Lett.*, 102:180504, 2009.

[132] Robert J. Collins, Ross J. Donaldson, Vedran Dunjko, Petros Wallden, Patrick J. Clarke, Erika Andersson, John Jeffers, and Gerald S. Buller. Realization of Quantum Digital Signatures without the Requirement of Quantum Memory. *Physical Review Letters*, 113(4):040502, July 2014.

[133] S. Ya. Kilin and A. B. Mikhalychev. Optical qudit-type entanglement creation at long distances by means of small cross-kerr nonlinearities. *Phys. Rev. A*, 83:052303, 2011.

[134] P. van Loock, T. D. Ladd, K. Sanaka, F. Yamaguchi, Kae Nemoto, W. J. Munro, and Y. Yamamoto. Hybrid quantum repeater using bright coherent light. *Phys. Rev. Lett.*, 96:240501, 2006.

[135] S. Olivares M. G. A. Paris B. Teklu, J. Trapani. Noisy quantum phase communication channels. *Physica Scripta*, 90:074027, 2013.

[136] Daiji Fukuda, Go Fujii, Takayuki Numata, Kuniaki Amemiya, Akio Yoshizawa, Hidemi Tsuchida, Hidetoshi Fujino, Hiroyuki Ishii, Taro Itatani, Shuichiro Inoue, and Tatsuya Zama. Titanium-based transition-edge photon number resolving detector with 98% detection efficiency with index-matched small-gap fiber coupling. *Optics Express*, 19(2):870–875, 2011.

[137] J.M. Bennett. *(Polarizers) Handbook of Optics Volume II*. McGraw-Hill, New York, US, 2 edition, 1995.

[138] M. Bilkis, M. Rosati, R. Morral Yepes, and J. Calsamiglia. Real-time calibration of coherent-state receivers: learning by trial and error. *arXiv.org*, quant-ph/arXiv:2001.10283, 2020.

[139] S. Brandsen, K. D. Stubbs, and H. D. Pfister. Reinforcement Learning with Neural Networks for Quantum Multiple Hypothesis Testing. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 1897–1902, June 2020. ISSN: 2157-8117.

[140] Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, November 1997.

[141] Ralf C. Staudemeyer and Eric Rothstein Morris. Understanding LSTM – a tutorial into Long Short-Term Memory Recurrent Neural Networks. *arXiv:1909.09586 [cs]*, September 2019. arXiv: 1909.09586.

[142] Oludare Isaac Abiodun, Aman Jantan, Abiodun Esther Omolara, Kemi Victoria Dada, Nachaat AbdElatif Mohamed, and Humaira Arshad. State-of-the-art in artificial neural network applications: A survey. *Heliyon*, 4(11), November 2018.

[143] Ahmed Tealab. Time series forecasting using artificial neural networks methodologies: A systematic review. *Future Computing and Informatics Journal*, 3(2):334–340, December 2018.

[144] Felix A. Gers, Jürgen Schmidhuber, and Fred Cummins. Learning to Forget: Continual Prediction with LSTM. *Neural Computation*, 12(10):2451–2471, October 2000.

[145] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

[146] François Chollet et al. Keras. https://keras.io, 2015.

[147] Katya Scheinberg Andrew R. Conn and Luis N. Vicente. *Introduction to Derivative-Free Optimization*. SIAM, Philadelphia, 2009.

[148] Darrell Whitley. A genetic algorithm tutorial. *Statistics and Computing*, 4(2):65–85, June 1994.

[149] Rainer Storn and Kenneth Price. Differential Evolution – A Simple and Efficient Heuristic for global Optimization over Continuous Spaces. *Journal of Global Optimization*, 11(4):341–359, December 1997.

[150] S. Das and P. N. Suganthan. Differential Evolution: A Survey of the State-of-the-Art. *IEEE Transactions on Evolutionary Computation*, 15(1):4–31, February 2011.

[151] Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. A Survey on Deep Transfer Learning. In Vera Kurková, Yannis Manolopoulos, Barbara Hammer, Lazaros Iliadis, and Ilias Maglogiannis, editors, *Artificial Neural Networks and Machine Learning – ICANN 2018*, Lecture Notes in Computer Science, pages 270–279, Cham, 2018. Springer International Publishing.

[152] S. J. Pan and Q. Yang. A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, October 2010.

[153] Alexander Semenovich Holevo. Bounds for the quantity of information transmitted by a quantum communication channel. *Problems of Information Transmission*, 9(3):177–183, 1973.

[154] A. S. Holevo and V. Giovannetti. Quantum channels and their entropic characteristics. *Reports on Progress in Physics*, 75(4):046001, March 2012.

[155] Jonathan L. Habif, Arunkumar Jagannathan, Samuel Gartenstein, Phoebe Amory, and Saikat Guha. Quantum-limited discrimination of laser light and thermal light. quant-ph/arXiv:1912.06718v1, 2019.

[156] A. G. Armada and M. Calvo. Phase noise and sub-carrier spacing effects on the performance of an ofdm communication system. *IEEE Communications Letters*, 2(1), 1998.

[157] Bing Qi, Lei-Lei Huang, Li Qian, and Hoi-Kwong Lo. Experimental study on the gaussian-modulated coherent-state quantum key distribution over standard telecommunication fibers. *Phys. Rev. A*, 76:052323, 2007.

[158] Paul Jouguet, Sébastien Kunz-Jacques, Anthony Leverrier, Philippe Grangier, and Eleni Diamanti. Experimental demonstration of long-distance continuous-variable quantum key distribution. *Nature Photonics*, 7:378, 2103.

[159] Duan Huang, Peng Huang, Dakai Lin, Chao Wang, and Guihua Zeng. High-speed continuous-variable quantum key distribution without sending a local oscillator. *Opt. Lett.*, 40(16):3695–3698, 2015.

[160] Adrien Marie and Romain Alléaume. Self-coherent phase reference sharing for continuous-variable quantum key distribution. *Phys. Rev. A*, 95:012316, 2017.

[161] J. R. Barry and J. M. Kahn. Carrier synchronization for homodyne and heterodyne detection of optical quadriphase-shift keying. *Journal of Lightwave Technology*, 10(12), 1992.

[162] D. Ly-Gagnon, S. Tsukamoto, K. Katoh, and K. Kikuchi. Coherent detection of optical quadrature phase-shift keying signals with carrier phase estimation. *Journal of Lightwave Technology*, 24(1):12–21, January 2006.

[163] Mohamed Morsy-Osman, Qunbi Zhuge, Lawrence R. Chen, and David V. Plant. Joint mitigation of laser phase noise and fiber nonlinearity for polarization-multiplexed qpsk and 16-qam coherent transmission systems. *Opt. Express*, 19(26):B329–B336, December 2011.

[164] Tao Wang, Peng Huang, Shitu Wang, and Guihua Zeng. Carrier-phase estimation for simultaneous quantum key distribution and classical communication using a real local oscillator. *Phys. Rev. A*, 99:022318, 2019.

[165] E. Ip and J. M. Kahn. Feedforward Carrier Recovery for Coherent Optical Communications. *Journal of Lightwave Technology*, 25(9):2675–2692, September 2007.

[166] Bing Qi, Pavel Lougovski, Raphael Pooser, Warren Grice, and Miljko Bobrek. Generating the Local Oscillator "Locally" in Continuous-Variable Quantum Key Distribution Based on Coherent Detection. *Physical Review X*, 5(4):041009, October 2015.

[167] Daniel B.?S. Soh, Constantin Brif, Patrick J. Coles, Norbert Lütkenhaus, Ryan M. Camacho, Junji Urayama, and Mohan Sarovar. Self-Referenced Continuous-Variable Quantum Key Distribution Protocol. *Physical Review X*, 5(4):041010, October 2015.

[168] Jun He, Robert A. Norwood, Maïté Brandt-Pearce, Ivan B. Djordjevic, Milorad Cvijetic, Suresh Subramaniam, Roland Himmelhuber, Carolyn Reynolds, Pierre Blanche, Brittany Lynn, and Nasser Peyghambarian. A survey on recent advances in optical communications. *Comp. and Elec. Eng.*, 40(1):216–240, 2014.

[169] Erik Agrell, Magnus Karlsson, A R Chraplyvy, David J Richardson, Peter M Krummrich, Peter Winzer, Kim Roberts, Johannes Karl Fischer, Seb J Savory, Benjamin J Eggleton, Marco Secondini, Frank R Kschischang, Andrew Lord, Josep Prat, Ioannis Tomkos, John E Bowers, Sudha Srinivasan, Maïté Brandt-Pearce, and Nicolas Gisin. Roadmap of optical communications. *Jour. of Opt.*, 18:063002, 2016.

[170] G. Xie, A. Dang, and H. Guo. Effects of Atmosphere Dominated Phase Fluctuation and Intensity Scintillation to DPSK System. *2011 IEEE International Conference on Communications (ICC)*, 2011.

[171] H. Ghozlan and G. Kramer. On Wiener phase noise channels at high Signal-to-Noise Ratio. In *2013 IEEE International Symposium on Information Theory*, pages 2279–2283, July 2013. ISSN: 2157-8117.

[172] Mohammad Reza Khanzadi. *Phase Noise in Communication Systems: Modeling, Compensation, and Performance Analysis*. phdthesis, Chalmers University of Technology, November 2015.

[173] Gilad Goldfarb and Guifang Li. BER estimation of QPSK homodyne detection with carrier phase estimation using digital signal processing. *Optics Express*, 14(18):8043–8053, September 2006.

[174] J. Salz. Modulation and detection for coherent lightwave communications. *IEEE Communications Magazine*, 24(6):38–49, June 1986.

[175] Vedran Dunjko and Hans J. Briegel. Machine learning & artificial intelligence in the quantum domain: a review of recent progress. *Reports on Progress in Physics*, 81(7):074001, 2018.

[176] Alexey A. Melnikov, Hendrik Poulsen Nautrup, Mario Krenn, Vedran Dunjko, Markus Tiersch, Anton Zeilinger, and Hans J. Briegel. Active learning machine learns to create new quantum experiments. *Proceedings of the National Academy of Sciences*, 115(6):1221–1226, 2018.

[177] Alexander Hentschel and Barry C. Sanders. Machine learning for precise quantum measurement. *Phys. Rev. Lett.*, 104:063603, February 2010.

[178] Alessandro Lumino, Emanuele Polino, Adil S. Rab, Giorgio Milani, Nicolò Spagnolo, Nathan Wiebe, and Fabio Sciarrino. Experimental phase estimation enhanced by machine learning. *Phys. Rev. Applied*, 10:044033, October 2018.

[179] Lukas J. Fiderer, Jonas Schuff, and Daniel Braun. Neural-Network Heuristics for Adaptive Bayesian Quantum Estimation. *arXiv:2003.02183 [quant-ph]*, March 2020. arXiv: 2003.02183.

[180] Valeria Cimini, Ilaria Gianani, Nicolò Spagnolo, Fabio Leccese, Fabio Sciarrino, and Marco Barbieri. Calibration of quantum sensors by neural networks. *Phys. Rev. Lett.*, 123:230502, 2019.

[181] Taira Giordani, Alessia Suprano, Emanuele Polino, Francesca Acanfora, Luca Innocenti, Alessandro Ferraro, Mauro Paternostro, Nicolò Spagnolo, and Fabio Sciarrino. Machine learning-based classification of vector vortex beams. *Phys. Rev. Lett.*, 124:160401, 2020.

[182] Gregory R. Steinbrecher, Jonathan P. Olson, Dirk Englund, and Jacques Carolan. Quantum optical neural networks. *npj Quantum Information*, 5(1):1–9, July 2019.

[183] Kerstin Beer, Dmytro Bondarenko, Terry Farrelly, Tobias J. Osborne, Robert Salzmann, Daniel Scheiermann, and Ramona Wolf. Training deep quantum neural networks. *Nat. Commun.*, 11(1):808, 2020.

[184] Julius Wallnöfer, Alexey A. Melnikov, Wolfgang Dür, and Hans J. Briegel. Machine Learning for Long-Distance Quantum Communication. *PRX Quantum*, 1(1):010301, September 2020.

[185] F. N. Khan, Q. Fan, C. Lu, and A. P. T. Lau. An optical communication's perspective on machine learning and its applications. *Journal of Lightwave Technology*, 37(2):493–516, 2019.

[186] Javier Mata, Ignacio de Miguel, Ramón J. Durán, Noemí Merayo, Sandeep Kumar Singh, Admela Jukan, and Mohit Chamania. Artificial intelligence (AI) methods in optical networks: A comprehensive survey. *Optical Switching and Networking*, 28:43–57, April 2018.

[187] Mingzhe Chen, Ursula Challita, Walid Saad, Changchuan Yin, and Mérouane Debbah. Artificial Neural Networks-Based Machine Learning for Wireless Networks: A Tutorial. *IEEE Communications Surveys Tutorials*, 21(4):3039–3071, 2019.

[188] Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Rev. Mod. Phys.*, 91:045002, 2019.

[189] Sanjaya Lohani, Erin M. Knutson, Matthew O'Donnell, Sean D. Huver, and Ryan T. Glasser. On the use of deep neural networks in optical communications. *Appl. Opt.*, 57(15):4180–4190, 2018.

[190] J. Thrane, J. Wass, M. Piels, J. C. M. Diniz, R. Jones, and D. Zibar. Machine learning techniques for optical performance monitoring from directly detected pdm-qam signals. *Journal of Lightwave Technology*, 35(4):868–875, 2017.

[191] Xiaoxia Wu, Jeffrey A. Jargon, Ronald A. Skoog, Loukas Paraschis, and Alan E. Willner. Applications of artificial neural networks in optical performance monitoring. *J. Lightwave Technol.*, 27(16):3580–3589, 2009.

[192] D. Zibar, M. Piels, R. Jones, and C. G. Schäeffer. Machine learning techniques in optical communication. *Journal of Lightwave Technology*, 34(6):1442–1452, 2016.

[193] Sanjaya Lohani, Erin M. Knutson, and Ryan T. Glasser. Generative machine learning for robust free-space communication. *Communications Physics*, 3(177):1–8, October 2020.

[194] B. Karanov, M. Chagnon, F. Thouin, T. A. Eriksson, H. Bülow, D. Lavery, P. Bayvel, and L. Schmalen. End-to-end deep learning of optical fiber communications. *Journal of Lightwave Technology*, 36(20):4843–4855, 2018.

[195] D. Zibar, L. H. H. de Carvalho, M. Piels, A. Doberstein, J. Diniz, B. Nebendahl, C. Franciscangelis, J. Estaran, H. Haisch, N. G. Gonzalez, J. C. R. F. de Oliveira, and I. T. Monroy. Application of machine learning techniques for amplitude and phase noise characterization. *Journal of Lightwave Technology*, 33(7):1333–1343, 2015.

[196] Faisal Nadeem Khan, Kangping Zhong, Xian Zhou, Waled Hussein Al-Arashi, Changyuan Yu, Chao Lu, and Alan Pak Tao Lau. Joint OSNR monitoring and modulation format identification in digital coherent receivers using deep neural networks. *Optics Express*, 25(15):17767–17776, 2017.

[197] D. Wang, M. Zhang, Z. Li, J. Li, M. Fu, Y. Cui, and X. Chen. Modulation format recognition and osnr estimation using cnn-based deep learning. *IEEE Photonics Technology Letters*, 29(19):1667–1670, 2017.

[198] Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. Rectifier nonlinearities improve neural network acoustic models. In *ICML*, volume 30, 2013.

[199] Itamar Holzman and Yachin Ivry. Superconducting nanowires for single-photon detection: Progress, challenges, and opportunities. *Advanced Quantum Technologies*, 2(3):1800058, 2019.

[200] T. Tieleman and G. Hinton. Lecture 6.5 - rmsprop, coursera: Neural networks for machine learning, 2012.

[201] G. E. Uhlenbeck and L. S. Ornstein. On the theory of the brownian motion. *Phys. Rev.*, 36:823–841, Sep 1930.

[202] Daniel T. Gillespie. The mathematics of brownian motion and johnson noise. *American Journal of Physics*, 64(3):225–240, 1996.

[203] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.

[204] Erich L. Lehmann and George Casella. *Theory of Point Estimation*. Springer-Verlag, New York, 2 edition, 1998.

[205] Luca Pezze' and Augusto Smerzi. Quantum theory of phase estimation. *arXiv:1411.5164 [cond-mat, physics:quant-ph]*, November 2014. arXiv: 1411.5164.

[206] Mankei Tsang, Howard M. Wiseman, and Carlton M. Caves. Fundamental quantum limit to waveform estimation. *Phys. Rev. Lett.*, 106:090401, 2011.

[207] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum metrology. *Phys. Rev. Lett.*, 96:010401, 2006.

[208] Marcin Jarzyna and Rafał Demkowicz-Dobrzański. Quantum interferometry with and without an external phase reference. *Phys. Rev. A*, 85:011801, January 2012.

[209] Mark Bradshaw, Ping Koy Lam, and Syed M. Assad. Ultimate precision of joint quadrature parameter estimation with a gaussian probe. *Phys. Rev. A*, 97:012106, January 2018.

[210] Emanuele Polino, Mauro Valeri, Nicolò Spagnolo, and Fabio Sciarrino. Photonic quantum metrology. quant-ph/arXiv:2003.05821v1, 2020.

[211] Dominik Šafránek. Simple expression for the quantum Fisher information matrix. *Physical Review A*, 97(4):042322, April 2018.

[212] Changhyoup Lee, Changhun Oh, Hyunseok Jeong, Carsten Rockstuhl, , and Su-Yong Lee. Using states with a large photon number variance to increase quantum fisher information in single-mode phase estimation. *arXiv*, page 190706349, 2019.

[213] Petr M. Anisimov, Gretchen M. Raterman, Aravind Chiruvelli, William N. Plick, Sean D. Huver, Hwang Lee, and Jonathan P. Dowling. Quantum metrology with two-mode squeezed vacuum: Parity detection beats the heisenberg limit. *Phys. Rev. Lett.*, 104:103602, March 2010.

[214] Kohjiro Iwasawa, Kenzo Makino, Hidehiro Yonezawa, Mankei Tsang, Aleksandar Davidovic, Elanor Huntington, and Akira Furusawa. Quantum-limited mirror-motion estimation. *Phys. Rev. Lett.*, 111:163602, 2013.

[215] Hidehiro Yonezawa, Daisuke Nakane, Trevor A. Wheatley, Kohjiro Iwasawa, Shuntaro Takeda, Hajime Arao, Kentaro Ohki, Koji Tsumura, Dominic W. Berry, Timothy C. Ralph, Howard M. Wiseman, Elanor H. Huntington, and Akira Furusawa. Quantum-enhanced optical-phase tracking. *Science*, 337(6101):1514–1517, 2012.

[216] Adriano A. Berni, Tobias Gehring, Bo M. Nielsen, Vitus Händchen, Matteo G. A. Paris, and Ulrik L. Andersen. Ab initio quantum-enhanced optical phase estimation using real-time feedback control. *Nature Photonics*, 9:577, 2015.

[217] U. Dorner, R. Demkowicz-Dobrzanski, B. J. Smith, J. S. Lundeen, W. Wasilewski, K. Banaszek, and I. A. Walmsley. Optimal quantum phase estimation. *Phys. Rev. Lett.*, 102:040403, January 2009.

[218] Nicolò Spagnolo, Chiara Vitelli, Vito Giovanni Lucivero, Vittorio Giovannetti, Lorenzo Maccone, and Fabio Sciarrino. Phase Estimation via Quantum Interferometry for Noisy Detectors. *Physical Review Letters*, 108(23):233602, June 2012.

[219] B. C. Sanders, G. J. Milburn, and Z. Zhang. Optimal quantum measurements for phase-shift estimation in optical interferometry. *Journal of Modern Optics*, 44(7):1309–1320, July 1997.

[220] Zhibo Hou, Rui-Jia Wang, Jun-Feng Tang, Haidong Yuan, Guo-Yong Xiang, Chuan-Feng Li, and Guang-Can Guo. Control-enhanced sequential scheme for general quantum parameter estimation at the heisenberg limit. *Phys. Rev. Lett.*, 123:040501, July 2019.

[221] Walker Larson and Bahaa E. A. Saleh. Supersensitive ancilla-based adaptive quantum phase estimation. *Phys. Rev. A*, 96:042110, October 2017.

[222] Kaimin Zheng, Huichao Xu, Aonan Zhang, Xinghai Ning, and Lijian Zhang. Ab initio phase estimation at the shot noise limit with on–off measurement. *Quantum Information Processing*, 18(11):329, September 2019.

[223] D. W. Berry, H. M. Wiseman, and J. K. Breslin. Optimal input states and feedback for interferometric phase estimation. *Phys. Rev. A*, 63:053804, April 2001.

[224] Zixin Huang, Keith R. Motes, Petr M. Anisimov, Jonathan P. Dowling, and Dominic W. Berry. Adaptive phase estimation with two-mode squeezed vacuum and parity measurement. *Phys. Rev. A*, 95:053837, May 2017.

[225] Brian E. Anderson, Prasoon Gupta, Bonnie L. Schmittberger, Travis Horrom, Carla Hermann-Avigliano, Kevin M. Jones, and Paul D. Lett. Phase sensing beyond the standard quantum limit with a variation on the su(1,1) interferometer. *Optica*, 4(7):752–756, July 2017.

[226] Brian E. Anderson, Bonnie L. Schmittberger, Prasoon Gupta, Kevin M. Jones, and Paul D. Lett. Optimal phase measurements with bright- and vacuum-seeded su(1,1) interferometers. *Phys. Rev. A*, 95:063843, June 2017.

[227] Prasoon Gupta, Bonnie L. Schmittberger, Brian E. Anderson, Kevin M. Jones, and Paul D. Lett. Optimized phase sensing in a truncated su(1,1) interferometer. *Opt. Express*, 26(1):391–401, January 2018.

[228] Shuro Izumi, Masahiro Takeoka, Kentaro Wakui, Mikio Fujiwara, Kazuhiro Ema, and Masahide Sasaki. Optical phase estimation via the coherent state and displaced-photon counting. *Phys. Rev. A*, 94:033842, 2016.

[229] Sergei Slussarenko, Morgan M. Weston, Helen M. Chrzanowski, Lynden K. Shalm, Varun B. Verma, Sae Woo Nam, and Geoff J. Pryde. Unconditional violation of the shot noise limit in photonic quantum metrology. *Nat. Photonics*, 11:700, 2017.

[230] Shakib Daryanoosh, Sergei Slussarenko, Dominic W. Berry, Howard M. Wiseman, and Geoff J. Pryde. Experimental optical phase measurement approaching the exact heisenberg limit. *Nature Communications*, 9:4606, 2018.

[231] B. L. Higgins, D. W. Berry, S. D. Bartlett, H. M. Wiseman, and G. J. Pryde. Entanglement-free heisenberg-limited phase estimation. *Nature*, 450:393–396, 2007.

[232] M. R. Vanner, J. Hofer, G. D. Cole, and M. Aspelmeyer. Cooling-by-measurement and mechanical state tomography via pulsed optomechanics. *Nature Communications*, 4(1):2295, 2013.

[233] Signe Seidelin, Yann Le Coq, and Klaus Mølmer. Rapid cooling of a strain-coupled oscillator by an optical phase-shift measurement. *Phys. Rev. A*, 100:013828, 2019.

[234]

[235] Isabelle Bouchoule and Klaus Mølmer. Preparation of spin-squeezed atomic states by optical-phase-shift measurement. *Phys. Rev. A*, 66:043811, 2002.

[236] Mankei Tsang and Ranjith Nair. Fundamental quantum limits to waveform detection. *Phys. Rev. A*, 86:042115, 2012.

[237] Giacomo M. D'Ariano, Matteo G. A. Paris, and Raffaella Seno. Feedback-assisted homodyne detection of phase shifts. *Phys. Rev. A*, 54:4495–4504, November 1996.

[238] H. M. Wiseman, D. W. Berry, S. D. Bartlett, B. L. Higgins, and G. J. Pryde. Adaptive measurements in the optical quantum information laboratory. *IEEE Journal of Selected Topics in Quantum Electronics*, 15(6):1661–1672, November 2009.

[239] H. M. Wiseman. Adaptive phase measurements of optical modes: Going beyond the marginal q distribution. *Phys. Rev. Lett.*, 75:4587–4590, December 1995.

[240] H. M. Wiseman and R. B. Killip. Adaptive single-shot phase measurements: A semiclassical approach. *Phys. Rev. A*, 56:944–957, July 1997.

[241] Igor Bargatin. Mutual information-based approach to adaptive homodyne detection of quantum optical states. *Phys. Rev. A*, 72:022316, August 2005.

[242] Michael A. Armen, John K. Au, John K. Stockton, Andrew C. Doherty, and Hideo Mabuchi. Adaptive homodyne measurement of optical phase. *Phys. Rev. Lett.*, 89:133602, 2002.

[243] Maria Bondani Matteo Bina, Alessia Allevi and Stefano Olivares. Phase-reference monitoring in coherent-state discrimination assisted by a photon-number resolving detector. *Scientific Reports*, (6):26025, 2016.

[244] D. W. Berry and H. M. Wiseman. Optimal states and almost optimal adaptive measurements for quantum interferometry. *Phys. Rev. Lett.*, 85:5098–5101, 2000.

[245] Wojciech Rzadkowski and Rafał Demkowicz-Dobrzanski. Discrete-to-continuous transition in quantum phase estimation. *Phys. Rev. A*, 96:032319, 2017.

[246] Liam Paninski. Asymptotic theory of information-theoretic experimental design. *Neural Computation*, 17(7):1480–1507, 2005.

[247] S. Kullback and R. A. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, March 1951.

[248] Solomon Kullback. *Information Theory and Statistics*. Number 11. Dover Publications, first edition, 1968.

[249] D. W. Berry, B. L. Higgins, S. D. Bartlett, M. W. Mitchell, G. J. Pryde, and H. M. Wiseman. How to perform the most accurate possible phase measurements. *Phys. Rev. A*, 80:052114, November 2009.

[250] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, USA, 2006.

[251] Nicolas Brunel and Jean-Pierre Nadal. Mutual Information, Fisher Information, and Population Coding. *Neural Computation*, 10(7):1731–1757, October 1998.

[252] A. S. Holevo, M. Sohma, and O. Hirota. Capacity of quantum Gaussian channels. *Physical Review A*, 59(3):1820–1828, March 1999.

[253] A. S. Holevo and R. F. Werner. Evaluating capacities of bosonic Gaussian channels. *Physical Review A*, 63(3):032312, February 2001.

[254] Vittorio Giovannetti, Seth Lloyd, Lorenzo Maccone, and Peter W. Shor. Entanglement Assisted Capacity of the Broadband Lossy Channel. *Physical Review Letters*, 91(4):047901, July 2003.

[255] Mark Wilde. *Quantum Information Theory 2nd Ed.* Cambridge University Press, Cambridge, United Kingdom, 2 edition, 2017.

[256] Narayanan Rengaswamy, Kaushik P. Seshadreesan, Saikat Guha, and Henry D. Pfister. Quantum-Message-Passing Receiver for Quantum-Enhanced Classical Communications. *arXiv:2003.04356 [quant-ph]*, March 2020. arXiv: 2003.04356 version: 1.

[257] Joseph M Renes. Belief propagation decoding of quantum channels by passing quantum messages. *New Journal of Physics*, 19(7):072001, jul 2017.

[258] Conor Delaney, Kaushik P. Seshadreesan, Ian MacCormack, Alexey Galda, Saikat Guha, and Prineha Narang. Demonstration of quantum advantage by a joint detection receiver for optical communications using quantum belief propagation on a trapped-ion device. *arXiv:2102.13052 [physics, physics:quant-ph]*, February 2021. arXiv: 2102.13052.

[259] Jian Chen, Jonathan L. Habif, Zachary Dutton, Richard Lazarus, and Saikat Guha. Optical codeword demodulation with error rates below the standard quantum limit using a conditional nulling receiver. *Nature Photonics*, 6, 2012.

[260] C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, July 1948.

[261] Jaehak Lee, Se-Wan Ji, Jiyong Park, and Hyunchul Nha. Classical capacity of Gaussian communication under a single noisy channel. *Physical Review A*, 91(4):042336, April 2015.

[262] Jaehak Lee, Se-Wan Ji, Jiyong Park, and Hyunchul Nha. Gaussian benchmark for optical communication aiming towards ultimate capacity. *Phys. Rev. A*, 93:050302, 2016.

[263] Robert G. W. Brown, Kevin D. Ridley, and John G. Rarity. Characterization of silicon avalanche photodiodes for photon correlation measurements. 1: Passive quenching. *Applied Optics*, 25(22):4122–4126, November 1986.

[264] Abdul Waris Ziarkash, Siddarth Koduru Joshi, Mario Stipcevic, and Rupert Ursin. Comparative study of afterpulsing behavior and models in single photon counting avalanche photo diode detectors. *Scientific Reports*, 8(1):5076, March 2018.