This is a repository copy of *Ultimate Limits of Thermal Pattern Recognition*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/184032/>

Version: Accepted Version

Article:

Harney, Cillian, Banchi, Leonardo and Pirandola, Stefano orcid.org/0000-0001-6165-5615 (2021) Ultimate Limits of Thermal Pattern Recognition. Physical Review A. 052406. ISSN 1094-1622

<https://doi.org/10.1103/PhysRevA.103.052406>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Ultimate Limits of Thermal Pattern Recognition

Cillian Harney¹, Leonardo Banchi^{2,3}, and Stefano Pirandola¹

¹*Department of Computer Science, University of York, York YO10 5GH, UK*

²*Department of Physics and Astronomy, University of Florence,
via G. Sansone 1, I-50019 Sesto Fiorentino (FI), Italy and*

³*INFN Sezione di Firenze, via G. Sansone 1, I-50019, Sesto Fiorentino (FI), Italy*

Quantum Channel Discrimination (QCD) presents a fundamental task in quantum information theory, with critical applications in quantum reading, illumination, data-readout and more. The extension to multiple quantum channel discrimination has seen a recent focus to characterise potential quantum advantage associated with quantum-enhanced discriminatory protocols. In this paper, we study thermal imaging as an environment localisation task, in which thermal images are modelled as ensembles of Gaussian phase insensitive channels with identical transmissivity, and pixels possess properties according to background (cold) or target (warm) thermal channels. Via the teleportation stretching of adaptive quantum protocols, we derive ultimate limits on the precision of pattern classification of abstract, binary thermal image spaces, and show that quantum-enhanced strategies may be used to provide significant quantum advantage over known optimal classical strategies. The environmental conditions and necessary resources for which advantage may be obtained are studied and discussed. We then numerically investigate the use of quantum-enhanced statistical classifiers, where quantum sensors are used in conjunction with machine learning image classification methods. Proving definitive advantage in the low-loss regime, this work motivates the use of quantum-enhanced sources for short-range thermal imaging and detection techniques for future quantum technologies.

The core aspiration of developing future quantum technologies is to exploit the intrinsic features of quantum mechanics in order to outperform any known optimal classical strategy for a given problem. Amidst rapidly accelerating progress in a number of fields of quantum computing [1–3], communications [4] and more, quantum sensing stands as the most mature and deployable of these fields [5]. This can be attributed, in large part, to a collection of theoretical advances in quantum decision theory/hypothesis testing [6–10], quantum metrology [11], and a focus on Continuous Variable (CV) bosonic quantum information which is particularly apt for quantum sensing [12–14].

An important setting of quantum hypothesis testing for a number of quantum technologies is Quantum Channel Discrimination (QCD), in which a user is tasked with classifying between a number of quantum channels using a (possibly) quantum-enhanced discrimination protocol. QCD finds crucial applications within quantum illumination [15–27], quantum reading [28–34]. Here, major quantum advantage has been theorised and in some cases experimentally verified [35–38], by exploiting non-classical properties of quantum input states and measurements.

Yet in general, this is a very difficult double optimisation problem, as one must determine the optimal input (or probe) state and optimal measurement used to minimise the error probability of misclassification. Furthermore, the most general protocol will make use of adaptive operations, which are extremely difficult to characterise and optimise. Nonetheless, critical performative insight into the fundamental lower bounds of discrimination error probabilities have been unveiled through the employment of channel simulation [39–41]. By simulating

channels through the use of quantum programmable processors, one can reduce channel discrimination to state discrimination and allow for the derivation of ultimate lower bounds. These methods are extremely powerful and have been used in a plethora of open problems in quantum communications, sensing and machine learning [42–46].

While the majority of QCD research has focussed on binary discrimination tasks, recent advances in the realm of multi-channel discrimination via the formulation of Channel Position Finding (CPF) [47, 48] has invited new exploration into the multi-ary domain. In this context, one may model images as consistent of pixels described by a target or a background quantum channel. This introduces a pattern recognition problem in which the goal is to perfectly classify a distribution of channels across an image.

Introduced in the context of barcode discrimination [49] in which pixels are modelled as pure lossy channels (relevant in the optical regime), these ideas can be extended into a thermal imaging setting. That is, thermal images may be described as a collection of pixels which are modelled by Gaussian phase insensitive channels subject to thermal noise. Formulated recently by Pereira et al. [50] in the context of single CPF, discrimination then becomes an environment localisation task in which one must discriminate between target/background channels of identical loss but different noise properties. Importantly, these collections of channels exhibit a property called joint teleportation-covariance, meaning one may employ multi-channel simulation techniques in order to derive ultimate error probability lower bounds. This presents a fascinating opportunity to unveil the optimal performance of quantum-enhanced thermal imag-

ing, and provide insight into its future applications and limitations.

This paper proceeds as follows: In Section I we review the image formulation introduced in [49, 50] in the context of environment localisation and the application of teleportation stretching that allows for the determination of ultimate lower bounds. In Section II we begin with the fundamental binary discrimination problem (a single pixel) in the context of thermal imaging, laying the foundations for exact pattern classification. Section III then makes use of statistical learning techniques in which a sensor is assisted by supervised learning, which allows us to explore larger image dimensions with lower resource demands.

I. TELEPORTATION STRETCHING & THERMAL IMAGING

A. Thermal Channel Patterns

Throughout this work we consider the following scenario: A thermal image $\mathbf{i}$ consists of m -pixels where each pixel is characterised by a Gaussian phase insensitive channel with identical transmissivities τ , but different environmental noise ν corresponding to a background channel $\mathcal{E}_B$ or a target channel $\mathcal{E}_T$. For instance, target channels may correspond to warmer pixels that we may wish to identify against a background of colder pixels. An m -pixel image can therefore be modelled as a channel pattern $\mathcal{E}_{\mathbf{i}}^m := \bigotimes_{k=1}^m \mathcal{E}_{i_k}$ where $i_k \in \{B, T\}$. Let

$$I_{[x]_{\mathbf{i}}}^m := \bigoplus_{k=1}^m x_{i_k} I, \quad (1)$$

be a function of a channel property x dependent on the position k in a pattern $\mathbf{i}$, where I is the 2×2 Identity matrix. Assuming an input m -mode Gaussian state fully characterised by its Covariance Matrix (CM) (with zero first moments), this multi-channel will perform the following transformation on its CM,

$$V_{\mathbf{i}} = (I_{[\sqrt{\tau}]_{\mathbf{i}}}^m) V (I_{[\sqrt{\tau}]_{\mathbf{i}}}^m)^T + I_{[\nu]_{\mathbf{i}}}. \quad (2)$$

Thermal-loss/amplifier channels are modelled by a beam-splitter of transmissivity τ , mixing the input signal with an environmental mode described by a thermal state of mean photon number $\bar{n}$. We consider there to exist background and target channels with equivalent transmissivities but different thermal numbers $\bar{n}_j$ for $j \in \{B, T\}$. Thermal-loss (amplifier) channels are parameterised by $0 \leq \tau < 1$ ($\tau > 1$) and $\nu \geq \frac{1-\tau}{2}$ for shot-noise $= \frac{1}{2}$. The induced noise of the target and background channels are related to the local thermal numbers in each mode as

$$\nu_j = \epsilon_j |1 - \tau|, \text{ such that } \epsilon_j := \bar{n}_j + \frac{1}{2}. \quad (3)$$

Meanwhile, Gaussian additive-noise channels are parameterised by $\tau = 1$ and $\nu_j \geq 0$, i.e. they are ideal channels with zero-loss, but are still subject to non-zero noise.

An ensemble of possible patterns occurring with probability $\pi_{\mathbf{i}}$ can define an image space $\{\pi_{\mathbf{i}}, \mathcal{E}_{\mathbf{i}}^m\} = \mathcal{U}$. Let us define some key binary image spaces: The set of all 2^m possible uniform channel patterns, with no global restrictions, is given by $\mathcal{U}_m$, studied as barcodes in [49]. If we restrict the space to the set of all m -channel patterns containing k target channels, then we call this problem k -Channel Position Finding (k -CPF) [47] since we are trying to locate the positions of these k channels. This image space is denoted by $\mathcal{U}_{\text{CPF}}^k$. More generally, we can define some fixed set of target channel numbers $\mathbf{k}$ characterising bounded k -CPF, $\mathcal{U}_{\text{CPF}}^{\mathbf{k}}$, such that there is uncertainty over the exact number of target channels present. Note that

$$\mathcal{U}_{\text{CPF}}^k \subseteq \mathcal{U}_{\text{CPF}}^{\mathbf{k}} \subseteq \mathcal{U}_m, \quad (4)$$

defines a hierarchy of image spaces, and the larger the image space, the more difficult discrimination becomes.

B. Multi-Channel Discrimination Protocols

The most general adaptive discrimination protocol is best described using a quantum comb [44, 51] (see Fig. 1(a)). This is a quantum circuit board that possesses M “slots” which may be filled with instances of m -length channel patterns $\mathcal{E}_{\mathbf{i}}^m$. The comb itself is a register of an arbitrary number of quantum systems, initially prepared in some state ρ_0 . Each channel pattern instance offers the opportunity to interrogate the multi-channel using some quantum systems within the register. Before and after each pattern instance, the discriminator may perform arbitrary joint quantum operations (QOs), which can be assumed to be trace-preserving [52]. This allows for the exploitation of unlimited entanglement, shared between input and output and feedback that can be used to adaptively optimise subsequent pattern interactions. After M adaptive probings, the comb is in its final state $\rho_{\mathbf{i}}^M$, and is subject to an optimal, joint POVM $\{\Pi_{\mathbf{i}'}\}_{\mathbf{i}' \in \mathcal{U}}$ which is classically post-processed to infer the channel pattern with some probability of error. Let us denote general adaptive protocols by $\mathcal{P}$, with classification error probability

$$p_{\text{err}}(\mathcal{P}) := \sum_{\mathbf{i} \neq \mathbf{i}'} \pi_{\mathbf{i}} \text{Tr} [\Pi_{\mathbf{i}'} \rho_{\mathbf{i}}^M]. \quad (5)$$

where $\sum_{\mathbf{i} \neq \mathbf{i}'}$ is a sum over all unequal channel patterns in the image space.

We may instead specify our strategy and consider block/block-assisted protocols, $\mathcal{B} \subset \mathcal{P}$. Here, a channel pattern $\mathcal{E}_{\mathbf{i}}^m$ is probed identically and independently by M -copies of an input state $\rho^{\otimes M} \rightarrow \rho_{\mathbf{i}}^{\otimes M} := \mathcal{E}_{\mathbf{i}}^m(\rho)^{\otimes M}$, followed by a generally joint POVM used to infer the pat-

Figure 1. Panels (a)-(c) depict the teleportation stretching of a general adaptive protocol, described by a quantum comb $\mathcal{P}$ that contains teleportation covariant channels. Via teleportation stretching, the general adaptive protocol is stretched to (b) a block-assisted protocol with the generation of M -copy Choi states, followed by a pre-measurement adaptive operation, $(\mathcal{B}, \tilde{\Lambda})$. The performance of (a) and (b) can then be bounded by the block-assisted protocol (c) in the limit of infinitely squeezed TMSV states in a probe/idler configuration, which achieves the ultimate performance limit.

tern. The classification error probability is consequently,

$$p_{\text{err}}(\mathcal{B}) := \sum_{i \neq i'} \pi_i \text{Tr} [\Pi_{i'} \rho_i^{\otimes M}]. \quad (6)$$

This protocol may be assisted by ancillary quantum systems that are free from channel interaction (often referred to as reference or idler modes), allowing for the use of entangled probe states which can lead to critical performance enhancements. The optimal block protocol is determined via a double optimisation over the input state, and the discrimination POVM. It is clear that the general adaptive protocol inherits all simpler block protocols and therefore the optimal error rate $p_{\text{err}}^{\text{opt}}$ always satisfies,

$$p_{\text{err}}^{\text{opt}}(\mathcal{P}) := \inf_{\mathcal{P}} [p_{\text{err}}(\mathcal{P})] \leq \inf_{\mathcal{B}} [p_{\text{err}}(\mathcal{B})] =: p_{\text{err}}^{\text{opt}}(\mathcal{B}) \quad (7)$$

C. Teleportation Stretching

Restriction to the space of block/block-assisted protocols $\mathcal{B}$ simplifies the investigation of optimal error bounds for multi-channel discrimination, at the expense of losing generality. However, there is a prominent class of channels for which the optimal block-assisted and general adaptive strategies coincide: Teleportation-covariant channels. A quantum channel admits telecovariance if for any input state ρ and pair of teleportation unitaries $\{U, V\}$, it satisfies,

$$\mathcal{E}(U\rho U^\dagger) = V\mathcal{E}(\rho)V^\dagger. \quad (8)$$

This condition means that we can *simulate* the channel $\mathcal{E}$ via a programmable teleportation protocol $\mathcal{T}$ using its Choi-state $\rho_{\mathcal{E}}$ as the resource (or program) state,

$$\mathcal{E}(\rho) = \mathcal{T}(\rho \otimes \rho_{\mathcal{E}}). \quad (9)$$

The simulation of multi-channels relies on the joint-telecovariance of a collection of channels i.e. $\mathcal{E}_i^m$ can be

simulated using the same teleportation protocol, but with (potentially) m different resource states,

$$\rho_{\mathcal{E}_i^m} := \rho_{\mathcal{E}_{i_1}} \otimes \cdots \otimes \rho_{\mathcal{E}_{i_m}} = \bigotimes_{k=1}^m \rho_{\mathcal{E}_{i_k}}. \quad (10)$$

Teleportation simulation refers to just one programmable protocol that can be used for channel simulation [39, 46], but is particularly convenient for telecovariant channels since their program states take such an expedient form.

Bosonic quantum channels require a particularly careful treatment. In this case, the Braunstein-Kimble teleportation protocol is invoked and the teleportation unitaries take the form of phase-space displacement operators [42, 53]. Since the Choi-states of bosonic channels are energy unbounded, the simulation is asymptotic. Then these Choi states are actually asymptotic quasi-Choi matrices, defined as the sequence of finite-energy Choi approximations in the limit of infinite squeezing,

$$\rho_{\mathcal{E}} := \lim_{a \rightarrow \infty} \{\rho_{\mathcal{E}}^a\} = \lim_{a \rightarrow \infty} (\mathcal{E} \otimes \mathcal{I})(\Phi_a). \quad (11)$$

Here Φ_a is a finite-energy TMSV state with $a = \bar{n}_S + \frac{1}{2}$, $\bar{n}_S$ is the mean photon number of the state, and $\mathcal{I}$ is the identity channel. This asymptotic treatment extends to all possible functionals taken on the quasi-Choi matrices, i.e. any n -state functional f of asymptotic Choi-states is computed in the infinite squeezing limit over Choi-sequences $(\{\rho_{\mathcal{E}_1}^a\}, \dots, \{\rho_{\mathcal{E}_n}^a\})$,

$$f(\rho_{\mathcal{E}_1}, \dots, \rho_{\mathcal{E}_n}) = \lim_{a \rightarrow \infty} f(\rho_{\mathcal{E}_1}^a, \dots, \rho_{\mathcal{E}_n}^a). \quad (12)$$

which is implicitly utilised throughout this paper.

The utility of telecovariance in the context of discrimination will now become clear; given a quantum comb $\mathcal{P}$, one may replace the channel pattern slots with their respective multi-channel teleportation simulations $\mathcal{E}_i^m(\rho) = \mathcal{T}(\rho \otimes \rho_{\mathcal{E}_i^m})$ (see Fig. 1(a)). Since teleportation is an LOCC protocol it can be summarised via a QO which is inherited by the subsequent adaptive operation Λ_k , ($k = 1, \dots, M$) that takes place at each pattern instance. This allows for the M channel resource

states $\rho_{\mathcal{E}_i^m}^{\otimes M}$ to be *stretched* back in time, outside of the adaptive operations at each round of pattern interaction. Furthermore, with all resource states stretched outside of the adaptive operations, these can be collapsed into a single trace-preserving QO, $\tilde{\Lambda}$.

Remarkably, teleportation stretching simplifies the quantum comb into the generation of M -resource states, followed by $\tilde{\Lambda}$ and a collective POVM, Π_i . The output state prior to measurement can be expressed as,

$$\rho_i^M = \tilde{\Lambda}(\rho_{\mathcal{E}_i^m}^{\otimes M}). \quad (13)$$

In the case of finite-dimensional systems, thanks to Naimark's dilation theorem any POVM can be represented as a quantum channel followed by a projective measurement [54, 55]. This means that $\tilde{\Lambda}$ can be further absorbed into the joint POVM $\tilde{\Pi}_i = \Pi_i \circ \tilde{\Lambda}$, completing the reduction from general adaptive protocol, to a block-assisted protocol $\mathcal{P} \rightarrow \mathcal{B}$. Then, the error rate of the optimal general adaptive protocol is equal to the error rate of the optimal block assisted protocol $p_{\text{err}}^{\text{opt}}(\mathcal{P}) = p_{\text{err}}^{\text{opt}}(\mathcal{B}, \tilde{\Lambda}) = p_{\text{err}}^{\text{opt}}(\mathcal{B})$, where $(\mathcal{B}, \tilde{\Lambda})$ is a block assisted protocol aided by a pre-measurement adaptive operation. Indeed, the optimal probability of misclassification for a generic image space $\mathcal{U} = \{\pi_i, \mathcal{E}_i^m\}$ is given by

$$p_{\text{err}}^{\text{opt}}(\mathcal{P}) = p_{\text{err}}^{\text{opt}}(\mathcal{B}) = \inf_{\tilde{\Pi}_i} \left[\sum_{i \neq i'} \pi_i \text{Tr} \left[\tilde{\Pi}_{i'} \rho_{\mathcal{E}_i^m}^{\otimes M} \right] \right]. \quad (14)$$

Deriving a lower bound for $p_{\text{err}}^{\text{opt}}(\mathcal{B})$ immediately implies a fundamental limit.

The same simplification may not apply for infinite dimensional systems, thus we provide a more general treatment by exploiting properties of the fidelity. Crucially, since trace preserving operations $\tilde{\Lambda}$ cannot increase the distance between two quantum states, one can write the data-processing inequality,

$$F(\rho_i^M, \rho_{i'}^M) = F\left(\Lambda\left(\rho_{\mathcal{E}_i^m}^{\otimes M}\right), \Lambda\left(\rho_{\mathcal{E}_{i'}^m}^{\otimes M}\right)\right), \quad (15)$$

$$\geq F\left(\rho_{\mathcal{E}_i^m}^{\otimes M}, \rho_{\mathcal{E}_{i'}^m}^{\otimes M}\right), \quad (16)$$

where the Bures fidelity $F(\rho, \sigma) = \|\sqrt{\rho}\sqrt{\sigma}\|_1 = \text{Tr}[\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}]$ is computed asymptotically (as in Eq. (12)). It then follows that this lower bound must hold for any possible general adaptive protocol $\mathcal{P}$, and all $i, i' \in \mathcal{U}$,

$$F_{\text{opt}} := \inf_{\mathcal{P}} [F(\rho_i^M, \rho_{i'}^M)] \geq F^M(\rho_{\mathcal{E}_i^m}, \rho_{\mathcal{E}_{i'}^m}). \quad (17)$$

It clear that this fidelity lower bound is achievable. Consider the block-assisted protocol $\mathcal{B}$, shown in Fig. 1(c), that probes each channel $i_k \in i$ in the pattern via M -copy TMSV states $\Phi_a^{\otimes M}$, while retaining idler modes for each channel. In the limit of infinite squeezing ($a \rightarrow \infty$) this results in the output state $\rho_{\mathcal{E}_{i'}^m}^{\otimes M}$ prior to collective

measurement, satisfying equality in Eq. (17). Thus the ultimate performance is achieved via $\mathcal{B}$ so that $F_{\text{opt}} = F^M(\rho_{\mathcal{E}_i^m}, \rho_{\mathcal{E}_{i'}^m})$. Substituting this result into the fidelity based state lower bound from [56], we gather

$$p_{\text{err}}^{\text{opt}}(\mathcal{P}) \geq \frac{1}{2} \sum_{i \neq i'} \pi_i \pi_{i'} F^{2M}(\rho_{\mathcal{E}_i^m}, \rho_{\mathcal{E}_{i'}^m}). \quad (18)$$

This represents a readily computable ultimate lower bound on the optimal error probability of classifying binary, jointly telecovariant channel patterns.

D. Ultimate Limits and Quantum Advantage

In order to know when quantum protocols are advantageous over classical protocols, it is necessary to derive an upper bound on $p_{\text{err}}^{\text{opt}}(\mathcal{P})$ that encompasses the use of quantum resources. An expedient quantum upper bound can be derived via a restriction to local measurements on the optimal block-assisted protocol $\mathcal{B}^{\text{loc}}$, such that $\Pi_j = \bigotimes_{k=1}^m \Pi_{j_k}$. The discriminatory POVMs are now completely separable, and as such the global classification error is defined as the compounded error of individual channel discriminations, $p_{\text{err}}^{\text{opt}}(\mathcal{P}) \leq p_{\text{err}}^{\text{opt}}(\mathcal{B}^{\text{loc}}) = 1 - (1 - p_{\text{err}}^{\text{pixel}})^m$ achieved by local Helstrom measurements. Using the analyses from [28, 49], considering m -length, uniformly distributed channel patterns probed by M -copy infinitely squeezed TMSV states we can write

$$p_{\text{err}}^{\text{opt}}(\mathcal{P}) \leq 1 - \left[1 - \frac{1}{2} F^M(\rho_{\mathcal{E}_B}, \rho_{\mathcal{E}_T}) \right]^m. \quad (19)$$

For non-uniform channel distributions, it may instead be useful to utilise the upper bound,

$$p_{\text{err}}^{\text{opt}}(\mathcal{P}) \leq \sum_{i \neq i'} \sqrt{\pi_i \pi_{i'}} F^M(\rho_{\mathcal{E}_i^m}, \rho_{\mathcal{E}_{i'}^m}), \quad (20)$$

based on Pretty Good Measurements (PGMs) [57], however Eq. (19) is always superior for uniformly distributed patterns.

Furthermore, it is important to characterise a lower bound on the optimal performance of protocols limited to classical resources. The best known classical strategy is an unassisted block-protocol with $p_{\text{err}}^{\text{opt}}(\mathcal{B}^{\text{cl}})$ such that $\mathcal{B}^{\text{cl}} \subset \mathcal{B}$. Such protocols are restricted to the use of coherent, classical input states, but have access to global measurements. Fortunately, since the thermal images are modelled using phase-independent channels with identical transmissivities, the application of operations that induce displacement or phase shifts on the input states have no effect on the outputs. Also, the optimal classical input state is pure; hence, the optimal classical strategy may simply make use of M -copy vacuum states $|0\rangle^{\otimes M}$ in order to probe each channel, followed by a joint measurement. Substituting $\mathcal{E}_i^m(|0\rangle_1 \dots \otimes |0\rangle_m)^{\otimes M}$ into Eq. (18) the classical lower bound is derived.

It is now possible to qualify, and quantify quantum advantage. Guaranteed quantum advantage defines the point at which a quantum-enhanced protocol outperforms the optimal classical strategy for a given discrimination task. In the case of uniformly distributed channel patterns, using the knowledge that block-assisted protocols limited to local measurements upper bound the optimal quantum performance, we can say that quantum advantage is guaranteed when

$$p_{\text{err}}^{\text{opt}}(\mathcal{B}^{\text{loc}}) \leq p_{\text{err}}^{\text{opt}}(\mathcal{B}^{\text{cl}}). \quad (21)$$

Since we cannot work directly with these quantities, advantage is only guaranteed whenever the lower bound of the classical strategy is outperformed by the upper bound of the quantum-enhanced strategy. Defining upper and lower bounds for classical (quantum) discrimination protocols, $p_{\text{err}}^{\text{cl,U}}$ and $p_{\text{err}}^{\text{cl,L}}$ ($p_{\text{err}}^{\text{q,U}}$ and $p_{\text{err}}^{\text{q,L}}$) we define the *minimum guaranteed advantage* (MGA) as the minimum performance enhancement achieved by a quantum strategy over the optimal classical one,

$$\Delta p_{\text{err}}^{\text{min}} := p_{\text{err}}^{\text{cl,L}} - p_{\text{err}}^{\text{q,U}}. \quad (22)$$

If $\Delta p_{\text{err}}^{\text{min}} > 0$, then quantum advantage is guaranteed. One can also define the *maximum potential advantage* (MPA) as

$$\Delta p_{\text{err}}^{\text{max}} := p_{\text{err}}^{\text{cl,L}} - p_{\text{err}}^{\text{q,L}}. \quad (23)$$

This represents the maximum possible improvement quantum strategies can bring when the derived quantum lower bound is fundamental. It is also only a *potential* advantage, and $\Delta p_{\text{err}}^{\text{max}}$ may be positive when $\Delta p_{\text{err}}^{\text{min}} \leq 0$ in which case advantage is not certified.

When determining conditions for advantage, it is also useful to devise a quantity that generally measures resource demands. To this end, we define the quantity

$$\bar{M} := \frac{M}{m}, \quad (24)$$

as the *relative probe copy number* of a given block protocol. This describes the number of probe copies required for a given discrimination protocol *relative* to the dimension of the channel pattern in question. A large $\bar{M}$ occurs when $M \gg m$, and implies a high resource cost since each channel in the pattern requires many probings prior to measurement. Contrarily, a small value of $\bar{M}$ implies a very low resource cost relative to the size of the image. Determining a *minimum* relative probe copy number defining the minimum $\bar{M}$ required to guarantee quantum advantage,

$$\bar{M}_{\text{adv}} = \arg \min_{\bar{M}} (|\Delta p_{\text{err}}^{\text{min}}|), \quad (25)$$

i.e. $\bar{M}_{\text{adv}} = \bar{M}$ such that $\Delta p_{\text{err}}^{\text{min}} = 0$, which provides a very useful metric on the resource cost and feasibility of a discrimination protocol.

While this characterisation of advantage relies on the use of infinitely squeezed input states, realistic quantum protocols must make use of finite-energy inputs. Finite-squeezing effects can be taken into account via the consideration of finite-energy TMSV states as probe inputs, rendering output states to be Choi-state approximations. The introduction of any two-mode squeezing immediately introduces stronger distinguishability over the classical vacuum states, and considerable advantage can be attained for realistic energy resources (see Appendix B for more details).

II. EXACT CLASSIFICATION OF THERMAL IMAGES

A. Uniform Image Spaces

Consider the m -pixel uniform image space $\mathcal{U}_m$ such that each pixel is equally likely to be a target or background environment. Thanks to the multiplicativity of the fidelity, the multimodal fidelity of the multi-channel Choi states can be simplified to

$$F(\rho_{\mathcal{E}_i^m}, \rho_{\mathcal{E}_{i'}^m}) = F(\rho_{\mathcal{E}_B}, \rho_{\mathcal{E}_T})^{\text{hamming}(i, i')}. \quad (26)$$

Employing the teleportation stretching lower bound Eq. (18), and exploiting properties of the Hamming distance [49] one can then derive an ultimate lower bound in terms of fidelity between single copy Choi states. In conjunction with the $\mathcal{B}^{\text{loc}}$ upper bound Eq. (19), the optimal error probability is bounded via,

$$p_{\text{err}}^{\text{opt}} \geq \frac{[F^{2M}(\rho_{\mathcal{E}_T}, \rho_{\mathcal{E}_W}) + 1]^m - 1}{2^{m+1}}, \quad (27)$$

$$p_{\text{err}}^{\text{opt}} \leq 1 - \left[1 - \frac{1}{2}F^M(\rho_{\mathcal{E}_B}, \rho_{\mathcal{E}_T})\right]^m. \quad (28)$$

Using Bernoulli's inequality $(1+x)^n \geq 1+nx$ in conjunction with Eqs. (28) and (27) these bounds can be more succinctly presented,

$$\frac{m}{2^{m+1}}F^{2M} \leq p_{\text{err}}^{\text{opt}} \leq \frac{m}{2}F^M. \quad (29)$$

It is now easy to qualify and quantify guaranteed advantage for a uniformly distributed m -pixel classification problem,

$$\Delta p_{\text{err}}^{\text{min}} > 0 \implies F_{\text{cl}}^{2M} > 2^m F_{\text{q}}^M. \quad (30)$$

Using Eq. (30) and $M = \bar{M}m$, it can be shown that the minimum relative probe copy number required to guarantee quantum advantage is given by

$$\bar{M}_{\text{adv}} = \frac{\log 2}{2 \log(F_{\text{cl}}) - \log(F_{\text{q}})}, \quad (31)$$

Figure 2. Minimum relative probe copy number $\bar{M}_{\text{adv}}$ required for guaranteed quantum advantage, with respect to: (a) Thermal-loss channels under low-loss and different thermal noise parameters, and (b) additive noise channels under different noise conditions. In (a) we model images using thermal target/background channels with pixel temperature variance $\Delta^\circ\text{C}$ (in a microwave setting with wavelength $\sim 1\text{mm}$, see Appendix C). The darker plots describe images with this variance at higher temperatures within $[0, 50]^\circ\text{C}$, whilst the brighter plots describe inherently colder environments within $[-10, 40]^\circ\text{C}$. In (b), the background channel noise ν_B is varied against the fixed target noise for a number of examples.

which is analysed in Fig. 2. Thanks to the compactness of the additive-noise Choi state fidelities, we can write

$$\bar{M}_{\text{adv}} = \frac{\log 2}{\log(1 + \Delta_q) - 2\log(1 + \Delta_{\text{cl}})}, \quad (32)$$

$$\Delta_q := \frac{(\sqrt{\nu_B} - \sqrt{\nu_T})^2}{2\sqrt{\nu_B\nu_T}}, \quad (33)$$

$$\Delta_{\text{cl}} := \sqrt{(\nu_T + 1)(\nu_B + 1)} - \sqrt{\nu_B\nu_T} - 1, \quad (34)$$

while the thermal-loss/amplifier result is too long to display, but can of course be treated via Eq. (31). Using any $\bar{M} > \bar{M}_{\text{adv}}$ guarantees advantage.

Figure 3. Panels (a) and (b) consider ultimate bounds on thermal image classification based on a uniform distribution of m target channels, $\mathbf{i} \in \mathcal{U}_m$. Panel (a) considers $m = 9$ pixel images such that each channel is a thermal loss channel under conditions $\{\epsilon_B, \epsilon_T\} = \{18.5, 20.2\}$ and different values of transmissivity τ in the low-loss regime. In Panel (b) this is shown for additive-noise channel patterns $\nu_T = 0.01$ for $m = 9$ pixels and different values of background noisy channels $\nu_B > \nu_T$. Panels (c) and (d) focus on k -BCPF image spaces $\mathbf{i} \in \bigcup_{k \in \mathcal{K}} \mathcal{U}_{\text{CPF}}^k$. Panel (c) considers $m = 9$ pixel images such that each channel is a thermal loss channel under conditions $\{\tau, \epsilon_B, \epsilon_T\} = \{0.99, 18.5, 20.2\}$. In Panel (d) this is shown for additive-noise channel patterns $\{\nu_B, \nu_T\} = \{0.02, 0.01\}$ for $m = 50$ pixels.

Characterising $\bar{M}_{\text{adv}}$ allows us to predict the necessary resource demands of a given thermal imaging setting for general m and M . In Fig. 2(a) we report resource costs for pattern discrimination at different pixel temperature variances $\Delta^\circ\text{C}$. From these results, it can be seen that discrimination is (unsurprisingly) easier at lower temperatures, since colder temperatures will endue less decoherence onto incident probes. Furthermore, advantage has a much greater dependence on the pixel variance Δ , than the specific temperatures. Nonetheless, it is clear that quantum advantage is only viable in the very low-loss limit, since $\bar{M}_{\text{adv}} \rightarrow \infty$ for $\tau \sim 0.95$.

As $\tau \rightarrow 1$ the prospect of advantage becomes increasingly likely, and when additive noise channels are considered (such that $\tau = 1$), guaranteed advantage is readily attainable for a variety of channel pattern conditions. Advantage is achieved at a low resource cost when *either* the target or background channel are of low additive noise (seen in Fig. 2(b)). In some very low noise settings $\bar{M}_{\text{adv}} \approx 1$, meaning that the necessary number of probe copies to guarantee advantage scales with the dimension of the pattern $M \sim m$. If $\bar{M}_{\text{adv}} < 1$ then $M \lesssim m$ and discrimination is extremely cost effective via

quantum resources.

Further results are shown in Fig. 3 for $m = 9$ pixel images, where we plot the MGA and MPA Δp_{err} with respect to probe number, for both channels and channel properties. The initial point at which the MGA is plotted is the first value of M for which the advantage inequality is satisfied, and determines the minimum number of probes required to guarantee quantum advantage.

B. Channel Position Finding

As discussed in Section I, one may consider non-uniform image spaces such as those formulated using channel position finding. Obtaining *a priori* knowledge of the image space can lead to dramatic improvements in discrimination performance, and is immensely useful for quantum illumination, reading and more. We begin with the concept of k -CPF, such that each m -pixel image in the space $\mathbf{i} \in \mathcal{U}_{\text{CPF}}^k$ contains exactly k -target channels. Then the image space dimension has been dramatically reduced from 2^m to $|\mathcal{U}_{\text{CPF}}^k| = \binom{m}{k}$ (where $\binom{m}{k}$ is the binomial coefficient). Deriving error bounds requires us to find a solution to

$$D_m^k = \frac{1}{\binom{m}{k}} \sum_{(\mathbf{i} \neq \mathbf{i}') \in \mathcal{U}_{\text{CPF}}^k} f^{\text{hamming}(\mathbf{i}, \mathbf{i}')}.$$
 (35)

Once more exploiting properties of the Hamming distance and a straightforward counting argument (see Appendix E), we find a closed formula in terms of the standard hypergeometric function,

$$D_m^k[f] = {}_2F_1(-k, k-n, 1, f^2) - 1.$$
 (36)

Substituting into the bounds from Eqs. (18) and (20), and assuming each pattern in the image space occurs with equal probability, $\pi_{\mathbf{i}} = \binom{m}{k}^{-1}$, then

$$\frac{D_m^k[F^{2M}]}{2\binom{m}{k}} \leq p_{\text{err}}^{\text{opt}} \leq D_m^k[F^M].$$
 (37)

One may more generally consider the case that there exists an image space $\mathcal{U}_{\text{CPF}}^{\mathbf{k}} = \bigcup_{k \in \mathbf{k}} \mathcal{U}_{\text{CPF}}^k$ such that $\mathbf{k}$ contains all possible numbers of target channels in any image in the space i.e. we have both an upper and lower bound on the number of targets, $\mathbf{k} = \{k_{\min}, \dots, k_{\max}\}$, or even just a suspected set of T target numbers $\mathbf{k} = \{k_1, \dots, k_T\}$. We call this k -Bounded CPF (k -BCPF) as we bound the number of targets in any possible image, generalising the previous results. In order to derive error probability bounds we must solve,

$$D_m^{\mathbf{k}}(f) = \sum_{(\mathbf{i} \neq \mathbf{i}') \in \mathcal{U}_{\text{CPF}}^{\mathbf{k}}} f^{\text{hamming}(\mathbf{i}, \mathbf{i}')}.$$
 (38)

By splitting this quantity into its diagonal (patterns drawn from likewise image spaces) and off diagonal (patterns drawn from unequal image spaces) sums, then it

can be solved using similar counting arguments as before. Indeed, defining the following functionals,

$$D_m^{\mathbf{k}}(f) = \sum_{j \in \mathbf{k}} D_m^j(f) + \sum_{i \in \mathbf{k}} \sum_{l \neq i} \tilde{D}_m^{i,l}(f).$$
 (39)

where D_m^j is given in Eq. (36) and

$$\tilde{D}_m^{k,l}(f) = \binom{m}{l} \binom{l}{k} f^{k-l} {}_2F_1(-k, l-m, l-k+1, f^2).$$
 (40)

Then one can write upper and lower bounds for k -BCPF using Eqs. (36) and (39), in conjunction with uniform weights, $\tilde{\pi}_{\mathbf{k}} = [\sum_{i \in \mathbf{k}} \binom{m}{i}]^{-1}$,

$$\frac{\tilde{\pi}_{\mathbf{k}}^2}{2} D_m^{\mathbf{k}}[F^{2M}] \leq p_{\text{err}}^{\text{opt}} \leq \tilde{\pi}_{\mathbf{k}} D_m^{\mathbf{k}}[F^M].$$
 (41)

Fig. 3(c) and (d) depict examples of these bounds for $m = 9$ and $m = 50$ pixel images. These results emphasise that precise knowledge of a target channel distribution has a significant role on the amount of advantage attainable through quantum sources. Ambiguity over the target distribution (even by one pixel) can remove any enhancement obtained over the discrimination of uniform channel patterns.

C. Distinguishability and Decoherence

In all the settings studied thus far, it is immediately obvious that the number of probes required to prove advantage for thermal-loss channels is very high, and increasingly so for lower transmissivity. Protection of the shared probe-idler entanglement is of course critical to improve the distinguishability of output states. When considering additive-noise channels, we isolate the role of thermal noise on the ultimate performance limits. While noise on its own is capable of degrading the quality of this entanglement, at sufficiently low-levels the quantum advantage is expedient and very effective. But when loss is combined with this noise, this endues a rapid decay in the shared probe-idler entanglement, as these decohering effects cooperate to reduce the distinguishability of output states.

It is therefore clear that idler-assisted discrimination protocols will prove quantum advantage only in the low-loss limit. At high loss, the ultimate and classical lower bounds quickly coincide, removing any benefit from utilising quantum sources. Such properties can be satisfied in the short-range imaging of thermal environments.

III. QUANTUM-ENHANCED PATTERN RECOGNITION

When dealing with a small number of pixels, the bounds in the previous section emphasise that there exist quantum advantageous protocols for image classification,

achievable with realistic channel environments. However as the number of pixels increases, performing this exact classification method leads to a high resource demand in order to achieve low error probability, as the discriminator attempts to classify a single pattern from an exponentially growing set of possible patterns. Instead, we can make use of machine learning techniques, commonplace in image classification problems, and enhance them with quantum sensors.

A. Nearest Neighbour Classifier

Image classification can be reframed as follows: Suppose there exists a space of all possible m -pixel images, such that the probability of obtaining an image $\mathbf{i}$ is $p_{\mathbf{i}}$. Each image has an associated class (or label) c , connected via a conditional probability distribution $P(c|\mathbf{i})$. Since this probability distribution is generally unknown, we design a classifier $\tilde{c}(\mathbf{i})$ that approximates its behaviour with a probability of error E . The smaller the error probability, the closer the approximate distribution is to reality. That is, an optimal classifier satisfies,

$$\tilde{c}_{\text{B}}(\mathbf{i}) = \underset{c}{\operatorname{argmax}}[P(c|\mathbf{i})] \quad (42)$$

and is known as the Bayes classifier, such that

$$E_{\text{B}} = 1 - \mathbb{E}_{\mathbf{i}}[P(\tilde{c}_{\text{B}}(\mathbf{i})|\mathbf{i})] \quad (43)$$

is its associated Bayes risk, which represents the minimum possible error probability of such a classifier. If a classification problem involves images each of which belong to one of N distinct classes, then the $E_{\text{B}} = 0$, since there exists an optimal classifier that can perfectly distinguish between all images.

However as stated, $P(c|\mathbf{i})$ is generally unknown, and it is therefore necessary to provide an approximation to this decision rule. This is often achieved with the assistance of a training set $\mathcal{T} = \{c_k; \mathbf{i}_k\}_k^{\text{tr}}$ of size $T = |\mathcal{T}|$ which contains pre-labelled images. The decision rule can then be modelled with a conditional probability distribution $P(c|\mathbf{i}, \mathcal{T})$ such that the classification of an image $\mathbf{i}$ depends on expertise drawn from $\mathcal{T}$.

A well studied option is the Nearest Neighbour (NN) classifier, characterised by the function,

$$\tilde{c}_{\text{NN}}(\mathbf{i}) = c\left(\underset{\mathbf{i}' \in \mathcal{T}}{\operatorname{argmin}}[d(\mathbf{i}, \mathbf{i}')]\right) \quad (44)$$

such that $\tilde{c}_{\text{NN}}$ retrieves the label of the closest image $\mathbf{i}' \in \mathcal{T}$ to the sampled image $\mathbf{i}$ that we wish to classify. ‘‘Closeness’’ in this context is defined by some appropriate distance measure $d(\mathbf{i}, \mathbf{i}')$ natural to the data in question. Using the NN classifier, and given a training set $\mathcal{T}$ of size T , one can write the expected classification error as

$$E_{\text{NN}}^T = \mathbb{E}_{\mathbf{i}, \mathcal{T}} \left[\sum_{c \neq \tilde{c}_{\text{NN}}} P(c|\mathbf{i}) P(\tilde{c}_{\text{NN}}|\mathbf{i}, \mathcal{T}) \right] \quad (45)$$

where we have recognised that the dependence on $\mathcal{T}$ has made the classifier itself a random variable. This quantity is the expected probability of misclassification, taken over image space and training sets of dimension T .

It is now important to note that in the limit of $T \rightarrow \infty$, provided that all N classes are distinct and $N \ll 2^m$, the classifier tends towards optimality and the Bayes risk (as shown in [49]). As the training set increases in size, the likelihood of a sampled state $\mathbf{i}'$ that we wish to classify also being contained in $\mathcal{T}$ tends towards unity. That is the probability of locating the closest image $\mathbf{i}'$ to $\mathbf{i}$ becomes

$$\lim_{T \rightarrow \infty} P_{\text{NN}}(\mathbf{i}'|\mathbf{i}, \mathcal{T}) = \delta_{\mathbf{i}', \mathbf{i}} \quad (46)$$

and in turn,

$$\lim_{T \rightarrow \infty} E_{\text{NN}} = E_{\text{B}} = 0. \quad (47)$$

Hence, in the limit of infinite training set size the NN classifier achieves the Bayes rate. In reality we are unable to make use of infinite training sets, but it is still possible to numerically investigate finite-sample-risk and study the performance of such classifiers. To this end, one can make use of the result from [58],

$$E_{\text{NN}}^T \approx E_{\text{NN}}^{\infty} + \sum_{j=2}^{\infty} x_j T^{-j/m} \quad (48)$$

which we can treat as an ansatz to interpolate numerical results. Truncating this series at a sufficient value, the finite sample error rate can be readily approximated.

B. Convolutional Neural Networks

Given a set of training data as before $\mathcal{T} = \{c_k; \mathbf{i}_k\}_k^{\text{tr}}$, a neural network is a universal approximating function $f_{\boldsymbol{\theta}}$ dependent on a set of variational parameters $\boldsymbol{\theta}$ that we wish to optimise in order to *learn* a relationship between the input/outputs of the training set. The goal of training is to adjust $\boldsymbol{\theta}$ such that

$$f_{\boldsymbol{\theta}}(\mathbf{i}_k) \approx c_k, \forall k. \quad (49)$$

Using a sufficiently large training set and variational parameter set, the network should be able to unveil (potentially highly non-linear) relationships between input images and their corresponding classification.

Feedforward neural networks form a particular architecture for this variational function, consistent of sequential layers of transformations applied to the original input. For an M -layer network acting as a classifier, an initial input vector $\mathbf{i}^{(1)}$ undergoes a series of transformations such that the final layer encodes an approximation to the desired output of the network,

$$\mathbf{i}^{(1)} \rightarrow \mathbf{i}^{(2)} \rightarrow \dots \rightarrow \mathbf{i}^{(M)} \mapsto \mathcal{C}(\mathbf{i}^{(M)}) \approx c.$$

Figure 4. Image classification via supervised machine learning post processing. Panel (a) depicts a (potentially quantum) sensor probing a thermal image, which is then classified via a convolutional neural network (CNN) which has been trained on similar data. The result of the classification is encoded in the final layer of the network. Panel (b) illustrates the use of a nearest neighbour classifier, which infers the class of a detected image by locating its closest (in terms of some distance measure) image in a training set $\mathcal{T}$.

for some decoding function, $\mathcal{C}(\cdot)$. The j^{th} element of the input vector transformed up to the k^{th} layer is given by

$$\mathbf{i}_j^{(k)} = \phi \left(\sum_l W_{jl} \mathbf{i}_l^{(k-1)} + b_j \right) \quad (50)$$

where ϕ is a suitable non-linear function and W_{jl} and b_j are real, network parameters. Throughout this work we utilise the Rectified Linear Unit (ReLU) as our simple non-linear activation function, where $\phi_{\text{RL}}(x) = \max(0, x)$. Such neural networks are typically optimised via a gradient descent learning scheme of some cost function $L(\theta)$ by means of backpropagation [59], a fundamental learning mechanism that formulates the underlying efficiency and performance of neural networks. For classification problems, the cost function may take numerous forms, such as the mean classification error, or cross entropy loss etc. that we wish to minimise.

Convolutional neural networks (CNNs) are a branch of network architectures that make use of convolutional layers of neurons, in which the hidden units have *local receptive fields* which are tied across the input [60]. The goal is to not only reduce the number of network parameters, but more importantly to extract critical features of input space that can be reused elsewhere, allowing the network to learn key characteristics with translational invariance. Figure 4(a) depicts the generic structure of a CNN. The transformation according to a convolutional layer is described by the operation of parallel filters which can simultaneously learn unique features.

Convolutional layers are ordered in a hierarchical manner, such that subsequent layers learn increasingly complex properties. For this reason, CNNs are particularly prevalent in image processing; early layers are used to identify basic, broad characteristics of images, such as edges and frames, whilst subsequent layers are then able to learn more complex structures, such as faces, hair, and texture. Furthermore, CNNs are very effective at noise reduction, i.e. performing effective image classification even amongst noisy inputs. In fact, noise can be shown to improve CNN classifier resilience [61], allowing the network to better generalise to a wide array of input possibilities rather than overfit on precise image patterns. This makes them an ideal image processing tool

that could be enhanced by means of quantum sensors for thermal imaging tasks.

C. Quantum Enhancement

We are now interested in determining how discrimination errors according to classical/quantum sensors propagate through to classification error of channel patterns when the statistical learning techniques of the previous section are employed.

Here we consider a practical classification example using the MNIST data-set [62], a commonly used data-set of 28×28 pixel images of handwritten digits 0-9, used for the benchmarking of many machine learning methods and models. The typical data-set is grey-scale, so we polarise the images such that each pixel only embodies one of two environments $\nu_j = \{\nu_B, \nu_T\}$. Since the channel distribution π_i is unknown and too difficult to characterise, it proposes an ideal application of data-driven techniques. Whilst this represents a toy example, it still captures the intent of thermal imaging and pattern classification. The data-set consists of a maximum training set $\mathcal{T}_{\text{max}} = \{\mathbf{i}_k; c_k\}_k$, and an evaluation set $\mathcal{V} = \{\mathbf{i}_l; c_l\}_l$ where $|\mathcal{T}_{\text{max}}| = 6 \times 10^4$ and $|\mathcal{V}| = 1 \times 10^4$. Using these data-sets we may quantify realistic error-rates on classification using quantum vs. classical sensors on thermal images.

Bounding the single pixel error probabilities according to [28]

$$\frac{1 - \sqrt{1 - F^{2M}}}{2} \leq p_{\text{err}}^{\text{pixel}} \leq \frac{F^M}{2}, \quad (51)$$

where F is the fidelity between background/channel outputs as before, we may then study how these pixel errors propagate to global image classification in both quantum/classical cases. The error E^T is then approximated as the expected probability of misclassification (using the NN or CNN classifier) across the evaluation set $\mathcal{V}$, using a random training subset $\mathcal{T} \subseteq \mathcal{T}_{\text{max}}$ where $|\mathcal{T}| = T$,

$$E^T \approx \frac{1}{|\mathcal{V}|} \sum_{\mathbf{i} \in \mathcal{V}} \delta(\tilde{c}(\mathbf{i}, \mathcal{T}), c(\mathbf{i})). \quad (52)$$

Figure 5. Thermal image classification performance on a thermally rendered MNIST handwritten digit data-set, using a NN Classifier. Each pixel is modelled as a thermal-loss channel with $\{\epsilon_B, \epsilon_T\} = \{18.5, 20.2\}$. The error region for the quantum-enhanced classifier at $M = 5000$ is extremely narrow, and essentially takes the form of the lower bound for $M = 2500$. Dashed lines indicate lower bound of the correspondingly coloured error region.

Here, $\tilde{c}(\mathbf{i}, \mathcal{T})$ is the inferred label of $\mathbf{i}$ from the T -size training set NN/CNN classifier, $c(\mathbf{i})$ is the true label, and δ is a Kronecker delta function such that

$$\delta(\tilde{c}(\mathbf{i}, \mathcal{T}), c(\mathbf{i})) = \begin{cases} 1, & \text{if } \tilde{c}(\mathbf{i}, \mathcal{T}) = c(\mathbf{i}), \\ 0, & \text{otherwise.} \end{cases} \quad (53)$$

Performing this computation using sensors with different properties, we can approximate regions within which the optimal error rate exists for finite training sets.

Figures 5 and 6 describe the results of our numerical simulations, analysing the error rates for both the NN and CNN classifiers for the thermal-loss and additive noise channels. In Fig. 5 (a) and (b) the performance of classical and quantum sensors respectively are investigated for the NN classifier and thermal-loss channels; plotting regions within which the error rate may lie (for numerous probe copy numbers M) against the dimension of the training set. These error ranges were numerically deduced through an empirical finite-sample risk analysis on an evaluation set using the NN classifier, for different sizes of training set T and $\tau = 0.99$, carried out for an optimal classical sensor (a) and a quantum sensor (b). The behaviour with respect to training set dimension was then interpolated via Eq. (48) (truncated at $j = 5$). Importantly, in the regime of low-loss quantum sources offer a more confident error rate, such that the separation between the upper and lower bounds is much narrower. As one would expect, increasing the training set dimension leads to significantly lower error rates, as the classifier becomes increasingly more likely to find a neighbouring sample within its training set that will provide a correct classification.

However, we see from Fig. 6 (a), quantum advantage for thermal-loss channels is very fragile with respect to loss. Here we focus on the CNN classifier, which is superior in its classification power (however this behaviour is

Figure 6. Quantum-enhanced thermal imaging on the MNIST data set, using NN and CNN pattern classifiers, for (a) thermal-loss channel patterns with $\{\epsilon_B, \epsilon_T\} = \{18.5, 20.2\}$ and (b) Additive noise channels with $\{\nu_B, \nu_T\} = \{0.02, 0.01\}$. Panel (a) depicts quantum advantage bounds for thermal-loss channels for variable τ , using a CNN classifier. Change in colour of the regions for $M = 500$ is used to indicate a transition from potential advantage to guaranteed advantage at $\tau \sim 0.99$ using realistic quantum resources. Panel (b) illustrates the error-classification bounds for additive noise channels under variable probe copy number M , using both statistical classifier models. In both cases $T = 10000$, dashed lines indicate lower bounds, and performances have been averaged over a number of experiments.

synonymous with the NN classifier). We plot the regions

$$[\Delta E^{\min}, \Delta E^{\max}] := [E_{\text{cl}}^L - E_{\text{q}}^U, E_{\text{cl}}^L - E_{\text{q}}^L]$$

(for the respective probe numbers) which are analogous definitions of MGA and MPA from Eqs. (22)-(23). We observe that even for transmissivities as high as $\tau \sim 0.97$ and a sufficient number of probe-copies, the advantage is either negligible or insufficiently guaranteed; only as τ approaches unity does advantage become guaranteed in relevant magnitudes (irrespective of the classifier). These results motivate the use of quantum sources only in a very

low-loss regime (concluded previously), as the decoherence involved with such channels removes advantageous entanglement between idler and signal probe too quickly.

As has been observed in previous sections, we can then expect much greater advantage associated with additive noise channels. Indeed, Fig. 6(b) shows the average classification error for additive noise channels for both classifiers with comparable resources $T = 10000$. Now we observe a significant level of quantum advantage, for a very low number of probes.

It is interesting to note that while both the CNN and NN classifiers follow very similar behaviours with respect to channel properties, there are a few key differences. In particular, provided with sufficient resources the CNN will converge to a superior misclassification rate than the NN classifier, clearly visible in Fig. 6(b). The simplicity of the NN approach means that its performance is wholly dependent on the dimension of its training set, and the precision of the sensor. Meanwhile, the CNN is capable of learning more intricate and detailed image features, meaning that it can more efficiently exploit the information provided from its training set.

Conversely, when the pixel error rate is very low, the CNN will typically perform slightly worse. This can be seen in the error regions of the classical sensors in Fig. 6(b). If the pixel error rate is too high, the CNN will struggle to extract any consistent features from a sampled pattern, making classification very difficult. Therefore in a high noise regime, NN classifiers may display improved robustness.

IV. DISCUSSION

Following from [48–50] we have formulated a problem setting for thermal imaging wherein the classification of an image is equated to the multi-channel discrimination of bosonic Gaussian channels. By considering binary collections of equally transmissive thermal-loss channels and additive noise channels, we exploited their joint teleportation covariance and applied the technique of teleportation stretching to derive ultimate lower bounds on the error probability of classification. This allowed for the nu-

merical and analytical analysis of these bounds for a number of image spaces, (uniform and non-uniform) under the formulation of channel position finding. These results emphasised the underlying quantum advantage that can be achieved in the limit of low-loss thermal channel patterns. Furthermore, we considered quantum-enhanced statistical pattern recognition for thermal imaging, making use of relevant machine learning methods in order to investigate the potential enhancement offered by quantum sources in a broader image setting.

Indeed, the limitation of entanglement enhanced discrimination resides in the susceptibility of such states to decoherence, which is particularly prominent in the presence of both loss and noise. While the discrimination of pure-loss and additive-noise channels can be impressively enhanced by quantum sources, the combination of these effects can dramatically degrade any guaranteed advantage.

Nonetheless, low-loss is ensured when the targets are limited to a short-range, making it easier to preserve quantum correlations between signal and idler-modes. This regime is realistic and important in future sensing applications, such as in low energy, non-invasive biomedical scanning. Alternative applications exist in communications, where eavesdropper activity along low-loss, multimode communication lines can be identified via their environmental noise properties. Eavesdroppers can be discriminated using appropriate environment localisation and pattern recognition techniques, such as those explored in this paper. These fundamental limits affirm the benefits and prospects of short-range, quantum-enhanced thermal imaging.

ACKNOWLEDGMENTS

C.H acknowledges funding from the EPSRC via a Doctoral Training Partnership (EP/R513386/1). S.P acknowledges funding from the European Union’s Horizon 2020 Research and Innovation Action under grant agreement No. 862644 (Quantum readout techniques and technologies, QUARTET). L.B. acknowledges support by the program “Rita Levi Montalcini” for young researchers.

-
- [1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*, 10th ed. (Cambridge University Press, USA, 2011).
 - [2] J. Preskill, *Quantum* **2**, 79 (2018).
 - [3] F. Arute *et al.*, *Nature* **574**, 505 (2019).
 - [4] S. Pirandola *et al.*, *Advances in Optics and Photonics* **12**, 1012 (2020).
 - [5] C. L. Degen, F. Reinhard, and P. Cappellaro, *Rev. Mod. Phys.* **89**, 035002 (2017).
 - [6] C. W. Helstrom, *Journal of Statistical Physics* **1**, 231 (1969).
 - [7] J. Bae and L.-C. Kwek, *Journal of Physics A: Mathematical and Theoretical* **48**, 083001 (2015).
 - [8] J. A. Bergou, U. Herzog, and M. Hillery, in *Quantum State Estimation* (Springer Berlin Heidelberg, 2004) pp. 417–465.
 - [9] A. Chefes, *Contemporary Physics* **41**, 401 (2000).
 - [10] G. Wang and M. Ying, *Phys. Rev. A* **73**, 042301 (2006).
 - [11] V. Giovannetti, S. Lloyd, and L. Maccone, *Nature Photonics* **5**, 222–229 (2011).
 - [12] S. Pirandola, B. R. Bardhan, T. Gehring, C. Weedbrook, and S. Lloyd, *Nature Photonics* **12**, 724 (2018).
 - [13] C. Weedbrook, S. Pirandola, R. García-Patrón, N. J. Cerf, T. C. Ralph, J. H. Shapiro, and S. Lloyd, *Rev.*

- Mod. Phys. **84**, 621 (2012).
- [14] A. Serafini, *Quantum Continuous Variables: A Primer of Theoretical Methods* (CRC Press, Taylor & Francis Group, 2017).
 - [15] S. Lloyd, Science **321**, 1463 (2008).
 - [16] S.-H. Tan, B. I. Erkmen, V. Giovannetti, S. Guha, S. Lloyd, L. Maccone, S. Pirandola, and J. H. Shapiro, Phys. Rev. Lett. **101**, 253601 (2008).
 - [17] J. H. Shapiro and S. Lloyd, New Journal of Physics **11**, 063045 (2009).
 - [18] Z. Zhang, M. Tengner, T. Zhong, F. N. C. Wong, and J. H. Shapiro, Phys. Rev. Lett. **111**, 010501 (2013).
 - [19] Z. Zhang, S. Mouradian, F. N. C. Wong, and J. H. Shapiro, Phys. Rev. Lett. **114**, 110506 (2015).
 - [20] S. Barzanjeh, S. Guha, C. Weedbrook, D. Vitali, J. H. Shapiro, and S. Pirandola, Phys. Rev. Lett. **114**, 080503 (2015).
 - [21] M. Dall'Arno, A. Bisio, and G. M. D'Ariano, International Journal of Quantum Information **10**, 1241010 (2012).
 - [22] Q. Zhuang, Z. Zhang, and J. H. Shapiro, Journal of the Optical Society of America B **34**, 1567 (2017).
 - [23] Q. Zhuang, Z. Zhang, and J. H. Shapiro, Phys. Rev. A **96**, 020302(R) (2017).
 - [24] U. L. Heras, R. D. Candia, K. G. Fedorov, F. Deppe, M. Sanz, and E. Solano, Scientific Reports **7**, 9333 (2017).
 - [25] G. De Palma and J. Borregaard, Phys. Rev. A **98**, 012101 (2018).
 - [26] R. Nair and M. Gu, Optica **7**, 771 (2020).
 - [27] A. Karsa, G. Spedalieri, Q. Zhuang, and S. Pirandola, Phys. Rev. Research **2**, 023414 (2020).
 - [28] S. Pirandola, Phys. Rev. Lett. **106**, 090504 (2011).
 - [29] R. Nair, Phys. Rev. A **84**, 032312 (2011).
 - [30] J. P. Tej, A. R. U. Devi, and A. K. Rajagopal, Phys. Rev. A **87**, 052308 (2013).
 - [31] G. Spedalieri, C. Lupo, S. Mancini, S. L. Braunstein, and S. Pirandola, Phys. Rev. A **86**, 012315 (2012).
 - [32] G. Spedalieri, Entropy **17**, 2218 (2015).
 - [33] Q. Zhuang, Z. Zhang, and J. H. Shapiro, Phys. Rev. A **96**, 040304(R) (2017).
 - [34] O. Hirota, Quantum Measurements and Quantum Metrology **4**, 70 (2017).
 - [35] E. D. Lopaeva, I. Ruo Berchera, I. P. Degiovanni, S. Olivares, G. Brida, and M. Genovese, Phys. Rev. Lett. **110**, 153603 (2013).
 - [36] S. Barzanjeh, S. Pirandola, D. Vitali, and J. M. Fink, Science Advances **6**, eabb0451 (2020).
 - [37] G. Ortolano, E. Losero, S. Pirandola, M. Genovese, and I. Ruo-Berchera, Science Advances **7**, eabc7796 (2021).
 - [38] Y. Xia, W. Li, Q. Zhuang, and Z. Zhang, arXiv:2006.11962 (2020).
 - [39] S. Pirandola, R. Laurenza, C. Lupo, and J. L. Pereira, npj Quantum Information **5**, 50 (2019).
 - [40] S. Pirandola and C. Lupo, Phys. Rev. Lett. **118**, 100502 (2017).
 - [41] S. Pirandola, S. L. Braunstein, R. Laurenza, C. Ottaviani, T. P. W. Cope, G. Spedalieri, and L. Banchi, Quantum Science and Technology **3**, 035009 (2018).
 - [42] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, Nature Communications **8**, 15043 (2017).
 - [43] S. Tserkis, J. Dias, and T. C. Ralph, Phys. Rev. A **98**, 052335 (2018).
 - [44] R. Laurenza, C. Lupo, G. Spedalieri, S. L. Braunstein, and S. Pirandola, Quantum Measurements and Quantum Metrology **5**, 1 (2018).
 - [45] R. Laurenza, C. Lupo, S. Lloyd, and S. Pirandola, Phys. Rev. Research **2**, 023023 (2020).
 - [46] L. Banchi, J. Pereira, S. Lloyd, and S. Pirandola, npj Quantum Information **6**, 42 (2020).
 - [47] Q. Zhuang and S. Pirandola, Communications Physics **3**, 103 (2020).
 - [48] Q. Zhuang and S. Pirandola, Phys. Rev. Lett. **125**, 080505 (2020).
 - [49] L. Banchi, Q. Zhuang, and S. Pirandola, Phys. Rev. Applied **14**, 064026 (2020).
 - [50] J. L. Pereira, Q. Zhuang, and S. Pirandola, Phys. Rev. Research **2**, 043189 (2020).
 - [51] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Phys. Rev. Lett. **101**, 060401 (2008).
 - [52] Since we have an unlimited number of systems in the register, and thanks to the principle of deferred measurement.
 - [53] S. L. Braunstein and H. J. Kimble, Phys. Rev. Lett. **80**, 869 (1998).
 - [54] C. W. Helstrom, International Journal of Theoretical Physics **8**, 361 (1973).
 - [55] M. G. A. Paris, The European Physical Journal Special Topics **203**, 61 (2012).
 - [56] A. Montanaro, in *2008 IEEE Information Theory Workshop* (2008) pp. 378–380.
 - [57] H. Barnum and E. Knill, Journal of Mathematical Physics **43**, 2097 (2002).
 - [58] R. R. Snapp and S. S. Venkatesh, in *Proceedings of the 12th IAPR International Conference on Pattern Recognition, Vol. 3 - Conference C: Signal Processing (Cat. No. 94CH3440-5)*, Vol. 2 (1994) pp. 1–6 vol.2.
 - [59] M. A. Nielsen, *Neural networks and deep learning* (Determination Press, 2015).
 - [60] K. P. Murphy, *Machine learning: a probabilistic perspective* (MIT Press, Cambridge, MA, 2012).
 - [61] T. S. Nazaré, G. B. P. da Costa, W. A. Contato, and M. Ponti, in *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications*, edited by M. Mendoza and S. Velastín (Springer International Publishing, Cham, 2018) pp. 416–424.
 - [62] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, Proceedings of the IEEE **86**, 2278 (1998).
 - [63] L. Banchi, S. L. Braunstein, and S. Pirandola, Phys. Rev. Lett. **115**, 260501 (2015).

Appendix A: Probe State Fidelities

1. Choi State Fidelities

The covariance matrix of a finite-energy Choi state of a Gaussian phase insensitive channel is given by

$$V = \begin{pmatrix} aI & \sqrt{\tau(a^2 - 1/4)}Z \\ \sqrt{\tau(a^2 - 1/4)}Z & (a\tau + \nu)I \end{pmatrix}, \quad (\text{A1})$$

where $a = \bar{n}_S + 1/2$ is the squeezing parameter of Φ_a , defined by the average mean photon number $\bar{n}_S$ per mode. It is now possible to evaluate the fidelities between Choi states of these channels via multi-modal fidelity formulae from [63]. Consider the fidelity between

Choi-states of thermal-loss/amplifier channels with parameters $\{\tau, \epsilon_T, \epsilon_B\}$ where $\epsilon_j = \bar{n}_j + 1/2$ are the local mean photon numbers for $j = B, T$. Taken in the limit of infinite-squeezing ($a \rightarrow \infty$), the fidelity between the Choi matrices depends only on the thermal parameters of each channel and not on τ [50],

$$F_{\text{th}}(\epsilon_T, \epsilon_B) = \frac{\sqrt{\epsilon_T \epsilon_B + 1 + \sqrt{(4\epsilon_T^2 - 1)(4\epsilon_B^2 - 1)}}}{\sqrt{2}(\epsilon_T + \epsilon_B)}. \quad (\text{A2})$$

For Gaussian additive-noise channels with noise parameters $\{\nu_T, \nu_B\}$, the fidelity in the limit of infinite squeezing takes a more succinct form and we get

$$F_{\text{add}}(\nu_T, \nu_B) = \frac{2\sqrt{\nu_T \nu_B}}{\nu_T + \nu_B}. \quad (\text{A3})$$

2. Optimal Classical State Fidelities

The optimal classical strategy for environment localisation may simply make use of a block-protocol using M -copy vacuum states $|0\rangle^{\otimes M}$ in order to probe a channel pattern. For thermal-loss/amplifier channels, the optimal classical strategy provides an output fidelity [50],

$$F_{\text{th}}^{\text{cl}}(\tau, \epsilon_T, \epsilon_B) = \frac{\sqrt{\alpha + \delta} + \sqrt{\alpha - \beta}}{\beta}, \quad (\text{A4})$$

$$\alpha = 4\epsilon_T \epsilon_B |1 - \tau|^2 + 2(\epsilon_T + \epsilon_B)\tau |1 - \tau| + (1 + \tau^2),$$

$$\beta = 2(\tau + (\epsilon_T + \epsilon_B)|1 - \tau|).$$

For additive noise channels the optimal classical output fidelity is given by,

$$F_{\text{add}}^{\text{cl}} = \frac{1}{\sqrt{(\nu_T + 1)(\nu_B + 1)} - \sqrt{\nu_T \nu_B}} \quad (\text{A5})$$

These fidelities can then be used within error bounds to characterise the optimal classical performance.

Appendix B: Finite-Energy Considerations

In this work we showed that TMSV states in the limit of infinite-squeezing derive ultimate limits for multi-channel discrimination of telecovariant quantum channels. However, it is not immediately clear how realistic probe states (which necessarily undergo finite-squeezing) approach this ultimate limit, and what the energy requirements are to observe meaningful advantage. We can readily access this information by looking at the finite-energy formats of the Choi-state fidelities in Eqs. (A2)-(A3). By studying the behaviour of these fidelities with respect to increasing energy, one can observe the resources required to realistically approach the ultimate limits derived in the main text.

Figure 7. Finite-energy output fidelity behaviour for (a) thermal-loss channels with background/target parameters $\{\epsilon_B, \epsilon_T\} = \{18.5, 20.2\}$, and (b) additive-noise channels with target noise $\nu_T = 0.01$. Beginning in a vacuum state (optimal classical probe) the finite-energy Choi-state approximations quickly approach the optimal quantum probe state fidelities in the limit of infinite squeezing, for realistic energetic resources, shown here for $a \in \{2.5, 10, 100\}$.

Contrary to Eq. (A2), the finite-energy Choi-state fidelity associated with thermal-loss channels does depend on the transmissivity, given by

$$F_{\text{th}}^a(\tau, \epsilon_T, \epsilon_B) = \frac{\sqrt{2}(\sqrt{\xi + \omega} + \sqrt{\xi - \omega})}{\omega}, \quad (\text{B1})$$

$$\xi = 4\epsilon_T \epsilon_B + 4a^2(4\epsilon_T \epsilon_B + 1) + |1 - \tau|^2(4a^2 - 1) \quad (\text{B2})$$

$$\omega = \sqrt{(4\epsilon_T^2 - 1)(4\epsilon_B^2 - 1)} + 8a\tau|1 - \tau|(1 + \tau)^2(\epsilon_T + \epsilon_B),$$

$$\omega = 4(\tau + (\epsilon_T + 2a\epsilon_B)|1 - \tau|). \quad (\text{B3})$$

Meanwhile, the finite-energy Choi-state fidelity associated with probing additive-noise channels admits the following form,

$$F_{\text{add}}^a(\nu_T, \nu_B) = \frac{2a\sqrt{\nu_T \nu_B} + \sqrt{(2a\nu_T + 1)(2a\nu_B + 1)}}{2a(\nu_T + \nu_B) + 1}. \quad (\text{B4})$$

In both cases the optimal classical and quantum fidelities can be recovered by taking the correct energy limits respectively,

$$F_j = \lim_{a \rightarrow \infty} F_j^a, \quad F_j^{\text{cl}} = \lim_{a \rightarrow \frac{1}{2}} F_j^a, \quad (\text{B5})$$

where the subscript $j \in \{\text{th}, \text{add}\}$ is used to label either thermal-loss or additive-noise fidelities. Furthermore, recall that $a := \bar{n}_S + \frac{1}{2}$, where $\bar{n}_S$ is the mean photon energy of incident probe states.

Figure 7 illustrates how the finite-energy fidelities approach the ultimate limits. For thermal-loss channels of high-transmissivity in Fig. 7(a), we show that the fidelity gap is quickly closed by finitely-squeezed probe states, and that even a small amount of squeezing $a = 2.5$ ($\bar{n}_S = 2$) leads to a significant improvement over the classical output fidelity. For larger, but very much realistic squeezing $a = 100$, the ultimate limit is essentially saturated for $\tau \lesssim 0.99$.

A similar behaviour can be observed for additive-noise channels in Fig. 7(b), where the gap between the optimal classical and quantum output fidelities is much more pronounced. Nonetheless, this fidelity separation is rapidly traversed for realistic energetic resources, emphasising the feasibility of high performance, entangled probe state protocols.

Appendix C: Temperature of Thermal Environments

Throughout this paper we characterise thermal environments via their mean photon number using $\bar{n}_j/\epsilon_j$. Using a Boltzmann distribution these can be converted to temperatures that characterise background and target pixels of thermal images. Assuming that channels are probed using microwave radiation of wavelength $\lambda \sim 1\text{mm}$, then we find pixel temperatures (in Kelvin),

$$T_j = \frac{hc}{k\lambda \log\left(\frac{1}{\bar{n}_j} + 1\right)}. \quad (\text{C1})$$

where k is Boltzmann's constant, h is Planck's constant, and c is the speed of light.

Appendix D: Uniform Image Spaces

Probing each pixel with M copies of a chosen input state provides the ultimate lower bound (for uniform *a priori* probabilities) given by,

$$\frac{1}{2^{2m+1}} \sum_{\mathbf{i} \neq \mathbf{i}'} F^{2M}(\rho_{\mathbf{i}}, \rho_{\mathbf{i}'}) \leq p_{\text{err}} \leq \frac{1}{2^m} \sum_{\mathbf{i} \neq \mathbf{i}'} F^M(\rho_{\mathbf{i}}, \rho_{\mathbf{i}'}). \quad (\text{D1})$$

Since we are dealing with telecovariant channels, we fundamentally bound these error probabilities through teleportation stretching. Stretching the discrimination protocol by supplanting the respective output states with the Choi states of each channel, $\rho_{\mathcal{E}_i^m} = \bigotimes_{k=1}^m \rho_{\mathcal{E}_{i_k}}$ such that $i_k \in \{B, T\}$, the fidelity between any two distinct, unequal patterns will be given by

$$F(\rho_{\mathcal{E}_i}, \rho_{\mathcal{E}_{i'}}) = F(\rho_{\mathcal{E}_B}, \rho_{\mathcal{E}_T})^{\text{hamming}(\mathbf{i}, \mathbf{i}')} \quad (\text{D2})$$

Considering the quantity,

$$D_m(f) = \frac{1}{2^m} \sum_{\mathbf{i} \neq \mathbf{i}'} f^{\text{hamming}(\mathbf{i}, \mathbf{i}')} = (f+1)^m - 1, \quad (\text{D3})$$

we find a closed form for this quantity. This can be proven via recursion as in [49], or by rewriting $\sum_{\mathbf{i} \neq \mathbf{i}'} = \sum_{\mathbf{i}, \mathbf{i}'} - \sum_{\mathbf{i}=\mathbf{i}'}$ and using multiplicativity of the fidelity. Using $D_m(f)$ the ultimate lower bound in Eq. (27) can then be shown. An upper bound using Eq. (20) can also be derived in this manner, such that

$$p_{\text{err}} \leq [F^M(\rho_{\mathcal{E}_B}, \rho_{\mathcal{E}_T}) + 1]^m - 1, \quad (\text{D4})$$

however this was shown to always be looser than Eq. (19).

Appendix E: Channel Position Finding

1. k -CPF Bounds

Let $\mathcal{U}_{\text{CPF}}^k$ be the set of m pixel channel patterns with strictly k -targets, such that $|\mathcal{U}_{\text{CPF}}^k| = \binom{m}{k}$ (where $\binom{m}{k}$ is the binomial coefficient). We wish to compute the sum

$$D_m^k = \frac{1}{\binom{m}{k}} \sum_{(\mathbf{i} \neq \mathbf{i}') \in \mathcal{U}_{\text{CPF}}^k} f^{\text{hamming}(\mathbf{i}, \mathbf{i}')} \quad (\text{E1})$$

For $\mathbf{i}, \mathbf{i}' \in \mathcal{U}_{\text{CPF}}^k$ all possible Hamming distances take the values:

$$\text{hamming}(\mathbf{i}, \mathbf{i}') = 2(t-k), \quad \forall t \in \{k+1, \dots, 2k\}. \quad (\text{E2})$$

for maximum distance k (where all k target pixels are different) and minimum distance 2 (where only one pixel is different). When $k=1$ we reduce to CPF as usual.

Consider each $\mathbf{i} \neq \mathbf{i}'$; there exist $t \in \{k+1, \dots, 2k\}$ locations in total for us to insert our target channels, and $\binom{m}{t}$ ways for this to occur. We may then place k targets into one of the channel patterns $\mathbf{i}$, with $\binom{t}{k}$ ways to do this. Having used k of our t locations, there now only exist $k - (t-k)$ channels that can be freely placed into the second pattern $\mathbf{i}'$, with $\binom{k}{2k-t}$ ways to do this. Via this counting argument the above sum can be rephrased,

$$\begin{aligned} \sum_{(\mathbf{i} \neq \mathbf{i}') \in \mathcal{U}_{\text{CPF}}^k} f^{\text{hamming}(\mathbf{i}, \mathbf{i}')} &= \sum_{t=k+1}^{2k} \binom{m}{t} \binom{t}{k} \binom{k}{2k-t} f^{2(t-k)}, \\ &= \binom{m}{k} [{}_2F_1(-k, k-m, 1, f^2) - 1], \end{aligned} \quad (\text{E3})$$

where ${}_2F_1$ is the standard hypergeometric function. Now we can define optimal quantum upper bounds and ultimate lower bounds on thermal k -CPF pattern recognition. Given

$$D_m^k(f) = [{}_2F_1(-k, k-m, 1, f^2) - 1], \quad (\text{E4})$$

then the bounds, assuming uniform *a priori* probabilities $\pi_{\mathbf{i}} = \binom{m}{k}^{-1}$, can be written as,

$$p_{\text{err}}^{\text{opt}} \geq \frac{D_m^k[F^{2M}(\rho_{\mathcal{E}_T}, \rho_{\mathcal{E}_B})]}{2\binom{m}{k}}, \quad (\text{E5})$$

$$p_{\text{err}}^{\text{opt}} \leq D_m^k[F^M(\rho_{\mathcal{E}_T}, \rho_{\mathcal{E}_B})]. \quad (\text{E6})$$

2. k -Bounded CPF Bounds

If the prior distribution of target channels is not uniform, then we may consider k -BCPF, such that there exists an image space $\mathcal{U}_{\text{CPF}}^{\mathbf{k}} = \bigcup_{k \in \mathbf{k}} \mathcal{U}_{\text{CPF}}^k$ such that $\mathbf{k}$ contains all possible numbers of target channels in any image in the space. In order to derive error probability bounds we must solve,

$$D_m^{\mathbf{k}}(f) = \sum_{(\mathbf{i} \neq \mathbf{i}') \in \mathcal{U}_{\text{CPF}}^{\mathbf{k}}} f^{\text{hamming}(\mathbf{i}, \mathbf{i}')} \quad (\text{E7})$$

The property $\mathbf{i} \neq \mathbf{i}'$ only applies when each image is being drawn from the same set $\mathcal{U}_{\text{CPF}}^k$, otherwise they are unequal by construction, hence we define the following quantity for $k \neq l$

$$\tilde{D}_m^{k,l}(f) = \sum_{\mathbf{i} \in \mathcal{U}_{\text{CPF}}^k, \mathbf{i}' \in \mathcal{U}_{\text{CPF}}^l} f^{\text{hamming}(\mathbf{i}, \mathbf{i}')}, \quad (\text{E8})$$

which allows us to rewrite Eq. (E7) by splitting it into independent diagonal and off diagonal summations,

$$D_m^{\mathbf{k}}(f) = \sum_{j \in \mathbf{k}} D_m^j(f) + \sum_{i \in \mathbf{k}} \sum_{l \neq i} \tilde{D}_m^{i,l}(f). \quad (\text{E9})$$

These $\tilde{D}_m^{k,l}(f)$ represent off diagonal functionals, in the sense that they count contributions from dissimilar k -CPF sets. Given $\mathbf{i} \in \mathcal{U}_{\text{CPF}}^k$ and $\mathbf{i}' \in \mathcal{U}_{\text{CPF}}^l$ and assuming that $k < l$, all possible Hamming distances take the values:

$$\text{hamming}(\mathbf{i}, \mathbf{i}') = 2t - (k + l), \quad \forall t \in \{l, \dots, k + l\}, \quad (\text{E10})$$

where the minimal distance is $l - k$ (all of $\mathbf{i}'$'s targets align with $\mathbf{i}$'s) and maximal distance is $k + l$ (the least possible amount of alignment between the patterns). Applying an identical counting argument as for standard k -CPF, one can express the off diagonal sums as

$$\tilde{D}_m^{k,l}(f) = \sum_{t=l}^{k+l} \binom{m}{t} \binom{t}{l} \binom{l}{(k+l)-t} f^{2t-(k+l)}. \quad (\text{E11})$$

In order to retrieve the original case of $k = l$, we simply adjust the t domain to eliminate self similar terms i.e. the minimum value of t becomes $t = k + 1$. This is once again solved by the standard hypergeometric function as

$$\tilde{D}_m^{k,l}(f) = \binom{m}{l} \binom{l}{k} f^{k-l} {}_2F_1(-k, l - m, l - k + 1, f^2). \quad (\text{E12})$$

Using this result one can write optimal upper and lower bounds for k -BCPF using Eq's. E7 and E9, in conjunction with uniform weights, $\pi_{\mathbf{k}} = [\sum_{i \in \mathbf{k}} \binom{m}{i}]^{-1}$, then

$$p_{\text{err}}^{\text{opt}} \geq \frac{\pi_{\mathbf{k}}^2}{2} D_m^{\mathbf{k}} [F^{2M}(\rho_{\mathcal{E}_T}, \rho_{\mathcal{E}_B})], \quad (\text{E13})$$

$$p_{\text{err}}^{\text{opt}} \leq \pi_{\mathbf{k}} D_m^{\mathbf{k}} [F^M(\rho_{\mathcal{E}_T}, \rho_{\mathcal{E}_B})]. \quad (\text{E14})$$

These bounds represent (i) the ultimate lower bounds on thermal pattern classification *given* that the number of target channels is contained in $\mathbf{k}$, and (ii) a quantum upper bound on the optimal error probability based on this pattern assumption and joint measurements. When $|\mathbf{k}| = m$ (contains all possible target numbers) we reproduce bounds for uniform channel patterns, and when $|\mathbf{k}| = 1$ (contains only one target number) we reduce to k -CPF.