

Neural Network Phase Estimation in Photon-Starved Coherent Beam Combination

Maria Popławska^a, Antoni Mikos-Nuszkiewicz^{a b}, and Marcin Jarzyna^a

^aCentre for Quantum Optical Technologies, Centre of New Technologies, University of Warsaw, Banacha 2c, 02-097 Warsaw, Poland

^bFaculty of Physics, University of Warsaw, Pasteura 5, 02-093 Warsaw, Poland

ABSTRACT

Coherent beam combination is a scalable and efficient technique for the reception of photon-starved signals in free-space optical communication systems. Its performance relies on accurate real-time estimation of phase fluctuations induced by e.g. atmospheric turbulence. Such estimation must be computationally efficient to enable phase correction during signal acquisition. In this work, we propose a novel approach to phase estimation in cascade coherent beam combination based on recurrent neural networks. The proposed model employs Long Short-Term Memory units to exploit temporal correlations present in photon-counting measurements and to infer the corresponding phase corrections. The network approximates a Bayesian estimation procedure that provides statistically optimal phase estimates. The results demonstrate the potential of recurrent neural networks as an alternative to conventional estimation methods for real-time coherent beam combination in photon-starved optical communication systems.

Keywords: coherent beam combination, deep space optical communication, photon-starved communication, phase estimation, neural networks

1. INTRODUCTION

Reliable reception of weak optical signals is one of the main challenges in free-space optical communication. This problem is particularly important in deep-space links, where the received optical power is extremely low, placing the communication system in the photon-starved regime.¹ In such a case, efficient signal collection requires large receiving apertures. However, for ground-based receivers, increasing the aperture size raises the sensitivity to atmospheric turbulence, making adaptive optics necessary and significantly increasing the complexity and cost of the receiver system.

An alternative way of signal reception that avoids the problems associated with a large aperture size is the technique of coherent beam combination (CBC), which allows the combination of signal beams gathered by a number of small apertures.^{2–4} In the basic CBC configuration, two beams that are outputs of two small apertures interfere on a balanced beam splitter whose dark output port is monitored by an on-off direct detector. The atmospheric turbulence, among other effects, such as scintillation,⁵ causes the relative optical phase between the beams to fluctuate in time. As a consequence, an imperfect interference occurs resulting in photocounts observed in the dark port, which in turn allow to compensate for the phase mismatch using a feedback loop and a phase modulator. As the number of receiving apertures increases, CBC can be implemented in a hierarchical cascade architecture, where pairs of beams are combined successively over multiple stages.⁴ The advantage is that not only one can gather the signal from a large total area by expanding the array of small apertures, but also, by cascading the outputs and combining them in pairwise manner, one can obtain almost a double increase of the received signal intensity every cascade layer. Such an architecture provides a scalable approach to increasing the collected optical power and improving signal quality in the photon-starved regime.

The performance of both single-stage and cascade CBC systems critically depends on accurate and computationally efficient phase estimation for feedback control. In this work, we investigate the use of recurrent neural

Further author information: (Send correspondence to M.P.)

M.P.: E-mail: m.poplawska@cent.uw.edu.pl




Figure 1. Basic scheme for coherent beam-combination of two light beams. The optical phase of the incoming signal evolves according to a Wiener diffusion process, resulting in a relative phase difference ϕ between two beams of equal intensity I_0 collected by the receiving apertures. The beams interfere on a balanced beam splitter, whose dark output is monitored by a photon detector. The detection outcomes are processed by a phase estimator (PE), which determines the phase correction applied by the phase modulator (PM) in a feedback loop. The objective of the feedback is to maximize the intensity I_{bright} at the bright output (equivalently, to minimize the intensity I_{dark} at the dark output)

network (RNN) architecture based phase estimator. Our model employs Long Short-Term Memory (LSTM) units⁶ to capture temporal correlations in photon counting data. The network is trained and validated on simulated datasets using target phase estimates generated by a Bayesian estimation procedure,³ which provides optimal phase estimates.

2. COHERENT BEAM-COMBINATION

In the simplest binary coherent beam-combination module, seen in Fig. 1, two optical beams with a relative phase difference ϕ and equal intensity I_0 interfere on a 50:50 beam splitter. Due to the impact of e.g. the turbulent atmosphere, the phase fluctuates in time which is modeled by a Wiener diffusion process,^{3,4} i.e. during each time interval Δt , the phase increment $\Delta\phi$ is assumed to be an independent Gaussian random variable with zero mean and variance $2D\Delta t$, where D is the phase diffusion constant. Consequently, the phase evolves according to the discrete-time random walk

$$\phi(t + \Delta t) = \phi(t) + \Delta\phi, \quad \Delta\phi \sim \mathcal{N}(0, 2D\Delta t). \quad (1)$$

The dark port is monitored by a Geiger-mode photon detector, which forms a part of the feedback loop that estimates and applies a phase correction $\tilde{\phi}$ to one of the input beams before it reaches the beam splitter. The intensities at the output ports are given by

$$I_{\text{bright,dark}} = I_0 \left[1 \pm \cos(\phi - \tilde{\phi}) \right], \quad (2)$$

where bright and dark output ports are labeled due to the fact that for perfect phase correction output intensities are a result of constructive ($I_{\text{bright}} = 2I_0$) and destructive ($I_{\text{dark}} = 0$) interference, respectively.

In the photon starved regime one does not deal with continuous intensities but rather discrete photocounts, which, according to standard photodetection theory, are characterized by Poissonian statistics. The probabilities of observing a click ($k = 1$) or not ($k = 0$) during a measurement interval of duration Δt are respectively given by

$$p(k = 1|\phi) = 1 - \exp(-I_{\text{dark}}\Delta t), \quad p(k = 0|\phi) = \exp(-I_{\text{dark}}\Delta t). \quad (3)$$

Note that, as the difference between the control and true phases increases, the probability $p(k = 1|\phi)$ of observing a click also grows, meaning the detection events provide partial information about the unknown phase difference. Importantly, the above basic two-beam CBC modules can be joined together in a cascaded manner to provide a combination scheme for multiple input beams, as seen in Fig. 2(a).




Figure 2. (a) Two-stage cascade CBC scheme. In the first stage, light collected by four apertures is combined pairwise. The bright outputs are subsequently combined in the second stage. (b)–(d) Simulated examples of the true phase difference (black) and the phase estimates (blue, orange, and red) applied by the phase modulators at the CBC stages indicated by the color corresponding to (a). (e) Intensities measured at the bright outputs of the corresponding beam-combination stages. The beam-combination efficiencies, shown in parentheses, were calculated using Eq. (4) over 100000 time steps. Simulation parameters: $D = 0.05 \text{ rad}^2/\text{ms}$, $I_0 = 10 \text{ photons/ms}$, and $\Delta t = 0.05 \text{ ms}$.

The goal of the CBC is to minimize (maximize) the intensity in the dark (bright) port, which is equivalent to minimizing the number of photocounts observed in the dark port. This can be accomplished using various phase correction algorithms. In the Bayesian approach, the probability distribution of the phase is updated after each measurement according to Bayes' theorem, combining prior knowledge with the likelihood of the observed detection outcome.³ The resulting posterior distribution is used to obtain a maximum a posteriori (MAP) estimate of the phase correction $\tilde{\phi}$, which is subsequently applied through a feedback loop to maximize the power directed to the bright output port. Note that from the MAP estimate one does not know the appropriate sign of the phase correction, however, if the incorrect sign has been chosen, the probability of a click in the next time step grows according to Eq. 3 and another correction will be made in a short time. The effectiveness of the Bayesian update heuristic is partly due to this self-correcting behavior. Importantly, even though we assume throughout this work that the incoming beams have equal intensities, the analysis can be readily generalized to the more realistic case of beams with unequal intensities or in the presence of external noise or dark counts.⁴

In order to combine more beams, several basic two-beam schemes may be put in a cascaded network, as seen in Fig. 2(a). The first stage of the cascade combines pairs of input beams. Next the bright output ports of the first layer two-beam CBC schemes are combined in the second stage etc., until one is left with a single output beam, enabling a further increase in the signal intensity. Figs. 2(b)-(e) present the simulation results for such a two-stage cascaded CBC configuration with four input beams. It is seen in (Figs. 2(b),(c) and Fig. 2(d) respectively, that the phase correction follows the true value of phase both in two first-layer CBC schemes as well as in the single second layer setup. The corresponding intensities in the bright ports for each case are seen in Fig. 2(e).

The crucial parameter describing the performance of a CBC scheme is its efficiency, defined as the averaged ratio of the intensity leaving the scheme through the bright port to the total input intensity. For a cascaded scheme with l layers the efficiency is given by

$$\eta = \left\langle \frac{I_{\text{bright}}}{2^l I_0} \right\rangle \cdot 100\%, \quad (4)$$

where $\langle . \rangle$ denotes the average over many realizations of the correction protocol. It is seen in Fig. 2(e) that one can attain a high combining efficiency both in individual two-beam CBC schemes as well as in a multiple-beam cascaded architecture.




Figure 3. (a) A scheme of a neural-network based phase estimator. The input vector, containing the phase estimates and click history over the past N_s time steps, is fed into the neural network, whose architecture is schematically shown. The network output is a two-dimensional vector corresponding to the sine and cosine of the estimated phase correction to be applied by the phase modulator. (b) Two unit vectors representing true phase difference between the incoming beams (ϕ , yellow) and its estimate ($\tilde{\phi}$, red). Their dot product provides a measure of agreement between the predicted and true phases (cosine similarity), and the loss is defined as one minus this quantity.

3. NEURAL NETWORK APPROACH

The problem of phase estimation in photon-starved coherent beam combination can be formulated as a time-series prediction task, where the control phase must be updated sequentially after each detection event. Classical supervised learning approaches for such sequential problems include recurrent neural networks (RNNs), in particular with Long Short-Term Memory (LSTM) units,⁶ as well as transformer-based architectures based on the multi-head attention mechanism.⁷ The LSTMs in their original or modified versions are used for a wide range of tasks⁸ with dependencies of data superposed with noise such as voice⁹ and handwriting recognition,¹⁰ energy consumption¹¹ or financial forecasting.¹² In this work, we restrict ourselves to LSTM-based models due to their favorable trade-off between computational complexity and real-time sequence analysis performance. LSTM networks, originally introduced to mitigate the vanishing gradient problem in standard RNNs, are capable of learning temporal dependencies by maintaining a persistent hidden state. Their gating mechanism – consisting of input, output, and forget gates – allows the network to selectively retain or discard information over time.⁶

Because the phase to be estimated during CBC follows a diffusion process, consecutive measurements are statistically correlated rather than independent. Exploiting these temporal correlations is therefore essential for accurate phase estimation, motivating the use of recurrent neural networks such as LSTMs. While this work focuses on LSTM-based models, other machine learning approaches may also prove effective, including alternative data representations, different sequence modeling architectures, or reinforcement learning strategies for direct optimization of the phase control process.

3.1 The model

The objective of the proposed neural network is to approximate the Bayesian phase estimator while providing faster inference suitable for real-time feedback. The network operates in an online feedback configuration, where previous outputs are used to determine new control phases after the true phase difference between the beams has drifted and a new detection event occurs. Consequently, the network must infer the temporal context of previous correction attempts. In particular, it must distinguish between situations in which several phase corrections have already been applied unsuccessfully, resulting in additional detector clicks, and situations in which a detector click occurs after a long period without detection events. This temporal dependency further motivates the use of a recurrent architecture.

All results in the paper were obtained using the TensorFlow (Keras API) library for machine learning in Python.¹³ The proposed architecture, schematically illustrated in Fig. 3(a), consists of an input layer, an LSTM layer, a fully connected output layer, and a normalization layer. The input layer defines the expected input tensor




Figure 4. Two-stage cascade CBC scheme with neural-network based phase estimation. (a) Scheme of two-stage CBC. (b)–(d) Examples of the true phase difference (black) and the corresponding phase estimates (blue, orange, and red) applied by the phase modulators at each CBC stage. (e) Intensities measured at the bright outputs of the corresponding beam-combination stages. The beam-combination efficiencies, shown in parentheses, were calculated using Eq. (4) over 100000 time steps. Simulation parameters: $D = 0.05 \text{ rad}^2/\text{ms}$, $I_0 = 10 \text{ photons/ms}$, and $\Delta t = 0.05 \text{ ms}$.

shape as a sequence of length $N = 8$, where each time step is represented by a three-dimensional feature vector ($d = 3$). As the input layer serves only to specify the input format, it contains no trainable parameters. The three input channels, depicted in Fig. 3(a), correspond to: (i) the sine of the control phase applied at each time step of the input sequence, (ii) the corresponding cosine values, and (iii) a binary sequence indicating photon detection events, with values 0 and 1 denoting the absence or occurrence of a detection event, respectively.

The recurrent part of the network is implemented using a standard LSTM cell with a hidden-state dimension ($H=8$), employing the default sigmoid and hyperbolic tangent activation functions. The LSTM output is subsequently mapped by a fully connected (**Dense**) layer with two output neurons. A normalization layer projects the output onto the unit circle by applying L_2 normalization. These normalized outputs are interpreted as estimates of the sine and cosine of the phase correction (see Fig. 3(a)). The normalizing layer is optional or can be implemented outside the model. The loss function for a single prediction is defined as cosine distance on the unit circle:

$$\mathcal{L} = 1 - \cos(\tilde{\phi} - \tilde{\phi}_{NN}) \quad (5)$$

where the cosine of the phase difference is equivalent to the dot product calculated between the target phase correction ($\sin \tilde{\phi}, \cos \tilde{\phi}$) and the network's prediction of this correction ($\sin \tilde{\phi}_{NN}, \cos \tilde{\phi}_{NN}$), see Fig. 3(b). For small angular differences, the loss function satisfies $\mathcal{L} \approx (\tilde{\phi} - \tilde{\phi}_{NN})^2/2$ and therefore locally behaves as the mean squared error (MSE) function while naturally accounting for the periodicity of angular variables.

The generated dataset was randomly divided into training (80%) and validation (20%) subsets. The dataset consisted of sequences generated using Bayesian estimation method but modified such that after each timestep the likelihood was updated depending on the measurement outcome on the dark port but with phase corrections determined by MAP applied only if detector on the dark port clicked. During training, the samples in the training set were shuffled at the beginning of each epoch and processed in batches. For each batch, the loss was computed as the mean of the individual sample losses defined in Eq. 5, and the network parameters were updated using the Adaptive Moment Estimation (Adam) optimizer.¹³ The validation set was used to monitor the training process, and the learning rate was automatically reduced upon a plateau in the validation loss using the ReduceLROnPlateau¹³ scheduling strategy.

3.2 Results

The network architecture and hyperparameters were selected to minimize the inference time while maintaining the beam-combination efficiency within an arbitrarily chosen tolerance of 3% with respect to previously reported Bayesian estimation methods.^{3,4} Since the Bayesian estimator was used to generate the training targets, the efficiency of supervised learning neural-network based estimator is not expected to surpass its efficiency. Instead,

the goal of the proposed approach is to achieve comparable estimation accuracy with a different implementation of phase estimation to open a new way for possible reduction in phase inference time.

The minimal working neural network consisted of $H = 8$ hidden units and was obtained for input sequences consisting of past $N = 8$ steps. It is seen in Fig. 4 that neural network estimator enables efficient beam combination both in the basic two-beam scheme as well as in the cascaded scenario. Note, that in order to have a useful phase correction procedure the bandwidth of the feedback loop and time required for the phase estimation protocol need to be in principle smaller than the time interval between subsequent measurements and characteristic time scale in which diffusion process randomizes the phase between the beams too much. This depends on the hardware implementation and optimization of the runtime routine, however in the above case the simplest way in which it was achieved was to convert the network using `tensorflow.lite` and running it with `Interpreter` interface.

4. CONCLUSIONS

In this work, we demonstrated a neural-network-based approach for phase estimation in photon-starved coherent beam combination systems. A compact LSTM estimator was developed and evaluated using datasets generated by a Bayesian phase estimation procedure. By optimizing the network architecture and input history, we obtained a lightweight model capable of maintaining comparable beam-combination efficiency.

The presented results demonstrate the feasibility of using recurrent neural networks for low-latency phase estimation in CBC feedback loops. Although the neural estimator does not surpass the Bayesian method used for generating the training targets in terms of efficiency, it provides a different implementation strategy with potential advantages in inference speed and hardware deployment. Further optimization of the network architecture and inference implementation may enable real-time neural-network-based phase stabilization in scalable coherent beam combination systems.

ACKNOWLEDGMENTS

This work was supported by the project “Quantum Optical Technologies” (contract no. FENG.02.01-IP.05-0017/23) carried out within the Measure 2.1 International Research Agendas programme of the Foundation for Polish Science cofinanced by the European Union under the European Funds for Smart Economy 2021-2027 (FENG). This research was funded in part by the National Science Centre, Poland grant SONATA 19 No. 2023/51/D/ST7/02171.

REFERENCES

- [1] Jarzyna, M., Kunz, L., Zwoliński, W., Jachura, M., and Banaszek, K., “Photon information efficiency limits in deep-space optical communications,” *Optical Engineering* **63** (Jan. 2024).
- [2] Larsson, R., Schröder, J., Karlsson, M., and Andrekson, P. A., “Coherent combining of low-power optical signals based on optically amplified error feedback,” *Optics Express* **30**, 19441 (May 2022).
- [3] Mikos-Nuszkiewicz, A., Paczos, J., Banaszek, K., and Jarzyna, M., “Bayesian approach to coherent combination of single photon beams,” *Optics Express* **32**(16), 28769 (2024).
- [4] Mikos-Nuszkiewicz, A., Łukanowski, K., Banaszek, K., and Jarzyna, M., “Deep-space optical communication receiver based on single photon coherent beam combination,” (2025).
- [5] Lyras, N. K., Kourogiorgas, C. I., and Panagopoulos, A. D., “Scintillation due to atmospheric turbulence on optical geo and meo satellite links: Time series synthesizer for system level simulations,” *International Journal of Satellite Communications and Networking* **43**, 435–446 (Oct. 2020).
- [6] Hochreiter, S. and Schmidhuber, J., “Long short-term memory,” *Neural Computation* **9**, 1735–1780 (Nov. 1997).
- [7] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I., “Attention is all you need,” (2017).
- [8] Mahmoud, A. and Mohammed, A., [*A Survey on Deep Learning for Time-Series Forecasting*], 365–392, Springer International Publishing (Dec. 2020).

- [9] Graves, A., Mohamed, A.-r., and Hinton, G., “Speech recognition with deep recurrent neural networks,” in *[2013 IEEE International Conference on Acoustics, Speech and Signal Processing]*, 6645–6649, IEEE (May 2013).
- [10] Graves, A., Fernández, S., Liwicki, M., Bunke, H., and Schmidhuber, J., “Unconstrained on-line handwriting recognition with recurrent neural networks.,” **20** (01 2007).
- [11] Maity, A. and Tukurul, V., “Forecasting energy consumption using recurrent neural networks: A comparative analysis,” (2026).
- [12] Islam, M. and Hossain, E., “Foreign exchange currency rate prediction using a gru-lstm hybrid network,” *Soft Computing Letters* **3**, 100009 (Dec. 2021).
- [13] Abadi, M., Agarwal, A., Barham, P., et al., “Tensorflow: Large-scale machine learning on heterogeneous systems.” <https://www.tensorflow.org/> (2015). Software available from tensorflow.org.