OPEN ACCESS
IOP Publishing
Classical and Quantum Gravity
Class. Quantum Grav. **42** (2025) 045003 (12pp) https://doi.org/10.1088/1361-6382/ada864

Automated alignment of an optical cavity using machine learning

Jiayi Qin[1,*], Katherine Kinder[1], Shreejit Jadhav[2,3,4], Praneel Chugh[5] and Bram J J Slagmolen[1]

[1] OzGrav-ANU, Centre for Gravitational Astrophysics, College of Science, The Australian National University, Australian Capital Territory 2601, Australia
[2] Inter-University Centre for Astronomy and Astrophysics (IUCAA), Post Bag 4, Ganeshkhind, Pune 411 007, India
[3] Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn VIC 3122, Australia
[4] ARC Centre of Excellence for Gravitational Wave Discovery (OzGrav), Melbourne, Australia
[5] Monash University, Melbourne, Australia

E-mail: jiayi.qin@anu.edu.au

Received 22 September 2024; revised 12 November 2024
Accepted for publication 9 January 2025
Published 24 January 2025

Abstract

Optimised alignment is important in optical systems, particularly in high-precision instrumentation such as gravitational wave detectors, in order to maximise the sensitivity. During operations, high performing optical wavefront sensing and feedback systems are used to maintain optical cavity performance. However, the need for an automated initial alignment process arises after maintenance or large environmental disturbances such as earthquakes, as it can be challenging to manually achieve optimised as well as consistent optical alignments. In this study, a machine learning control system is presented to determine the optimal input beam alignment of an optical cavity based on a digital camera stream of the transmitted cavity mode. We use convolutional neural networks to classify the cavity mode from its image, with 100% prediction accuracy for the desired mode. A genetic algorithm is applied to find experimental parameters that maximise the transmitted power of a chosen cavity mode. The system demonstrates consistent alignment outcomes that the

* Author to whom any correspondence should be addressed.

Original Content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

median intensity over multiple trials exceeds 95% by the sixth generation of the algorithm. These results show that machine learning techniques can be implemented to automate the alignment process that is compatible for a broad range of optical resonator platforms.

Keywords: optical alignment, convolutional neural networks, genetic algorithm, reinforcement learning, gravitational wave detectors

1. Introduction

Many modern physics experiments use lasers in precision metrology [1–4], interactions with matter at the atomic scale [5], and revelation of fundamental cosmic information [6–8]. In an optical system, the cavity is a common element that is used to enhance the interaction between the optical fields and objects. Optical cavities achieve resonant fields through constructive interference and are widely used in optical control and sensing technologies [9, 10]. Gravitational-wave detectors, for example LIGO, Virgo and KAGRA [2, 11, 12], utilise optical cavities to increase the light interactions with the test masses for increased sensitivity. This makes it possible to detect the extremely weak ripples in space-time called gravitational waves that arise from energetic astrophysical events like collisions of compact objects, supernovae, perturbations in the early Universe, etc [13].

Optical cavities have various sizes and geometries, but all require the reflective components to be aligned in a designated configuration such that the light constructively interferes to create a resonant optical mode [14]. The most frequently used modes are fundamental Gaussian modes, while some optical experiments focus on alignment to specific higher-order modes [15, 16]. The alignments of the input beam and cavity optics are important to achieve maximal optical circulating powers as well as to mitigate optical instabilities caused by the resonance of higher-order optical modes in the cavity. Hence, techniques to optimise and maintain optical alignment are essential for the operation of precision instrumentation like gravitational wave detectors [17–19]. For example, during operation, gravitational wave detectors use angular control systems for automated optical alignments, like the differential wavefront sensing with quadrant photo-diodes [20–22] (GEO600) or the Anderson technique that introduces frequency side-bands [14, 23] (Virgo) to maximise the power in the arm cavity. However, these techniques often necessitate an initial optical alignment for the cavity to maintain resonance to acquire a control signal, a process that can be particularly challenging in large optical cavities with suspended mirrors. When many such suspended mirrors, like in gravitational wave detectors, get severely misaligned due to earthquakes or unintended maintenance events, the cavity resonance conditions will be lost. This often means that the detector needs to start its initial alignment procedure to re-align the mirrors such that they obtain optical resonance conditions, which can be time-consuming and temporarily take these detectors offline.

Machine learning algorithms offer a range of tools that can play a useful role in experimental research. Manual optimisations of experiments, like the alignment of optical cavities, often rely on human skill which may not offer the required accuracy and consistency for reaching the best solution [21, 24, 25]. In contrast, machine-based methods minimize the need for manual intervention, thereby saving time and reducing the potential for human error. Automation refines the alignment process, decreasing the frequency of necessary adjustments and ensuring more consistent and reliable outcomes. For example, aligning an optical cavity requires identification of the cavity mode, as well as strategies to find best parameters. For such tasks, machine learning algorithms, e.g. deep neural networks and reinforcement learning, can provide suitable solutions and automate experimental processes partly or in full [26–33]. Deep

neural networks are a class of machine learning models that can be particularly useful in the identification of resonant modes in a cavity. These models allow learning complex features through optimization of their parameters by training on large quantities of data [34, 35]. A specific class of neural networks called convolutional neural networks (CNNs) is especially useful for learning features from 2D data, like images [35–37]. The use of convolution allows CNNs to utilise the translation symmetry of features in images for efficient learning [35, 36, 38]. As the optical alignment presents a classic search problem, another helpful technique is reinforcement learning algorithm based on given reward functions [39]. The reward function is a numerical metric used to evaluate how well each candidate solution meets the optimization goal. Recently, reinforcement learning techniques have been successfully demonstrated for problems from alignment of a simple Mach-Zehnder interferometer [28] to complex sensing and control of a kilometer-scale gravitational wave detector [40]. A genetic algorithm (GA) is an optimisation strategy based on the principle of 'survival of the fittest' - a subset of top-performing solutions (parents) from a set of candidate solutions of a generation are used to create the next generation of solutions (offspring) until a termination condition is satisfied [41–43]. GAs with reinforcement learning have proven to have high accuracy and success in various control and optimisation applications [44–47].

In this paper, we present a machine-learning-based control scheme that combines a CNN with a GA to find the optimal alignment of a generic optical cavity. A stream of images capturing the optical mode at the cavity transmission is analysed by the CNN to construct a reward function which, in turn, is used by the GA to find the best alignment. This study demonstrates that such a scheme produces optimal and consistent alignments. Importantly, this scheme shows great simplicity in readout requirement—a single camera monitor, which is more versatile than traditional methods that need implementation of a photo-diode and camera. This technique can be applied to various systems of precision metrology that require initial optical cavity alignment, including the suspended optical cavities in gravitational-wave detectors. This proposed method enhances the alignment process by being effective even with less accurate initial setups. This is especially valuable for complex systems where achieving precise initial alignment is challenging. Designed to function effectively regardless of the starting alignment condition—whether well-aligned or completely misaligned—this method offers greater flexibility and ease of implementation.

2. Experimental setup

As shown in figure 1, a simple Fabry-Pérot cavity was built by using a planar input mirror and a concave output mirror which is attached to a piezo actuator (PZT) for cavity length control. Additional lenses are used to mode-match the laser beam into the cavity. To steer and control the cavity input beam, there are two magnet-coil actuated fast-steering mirrors (Newport FSM-300, resolution $\leqslant 1\,\mu$rad rms, accuracy$\leqslant$0.262 mrad rms) (FSM1 and FSM2) that provide four degrees of freedom (each mirror having pitch θ_x and yaw θ_y) [14]. A high-speed Gigabit-Ethernet camera (Basler acA720-290gm) is used to monitor and capture the transmitted optical cavity mode in real-time. The PZT, FSMs, and camera are connected to a computer system for control.

3. Machine learning algorithms

The machine learning algorithm in this study consists of two parts: a CNN, which identifies the cavity resonant mode by analysing the images collected at the cavity transmission; and a GA,

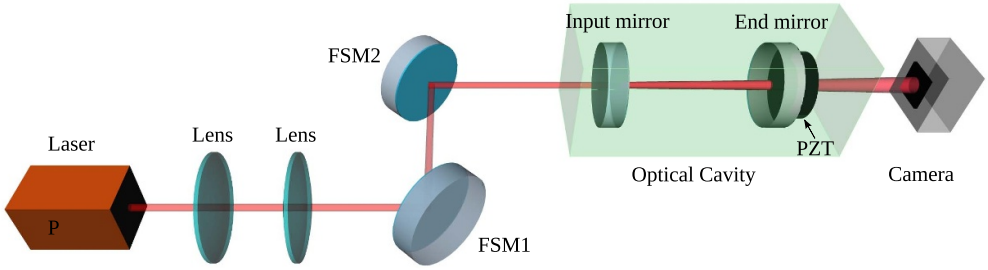

Figure 1. Experimental Setup. A simple Fabry-Pérot optical cavity is formed by a planar front mirror with a reflectivity of 97% and a concave end mirror with a reflectivity of 99%. Two lenses are used to mode match the waist of the incident Gaussian beam into that of the cavity. The laser wavelength is 1064 nm. The alignment of incident light is controlled by two fast-steering mirrors (FSM1 and FSM2).

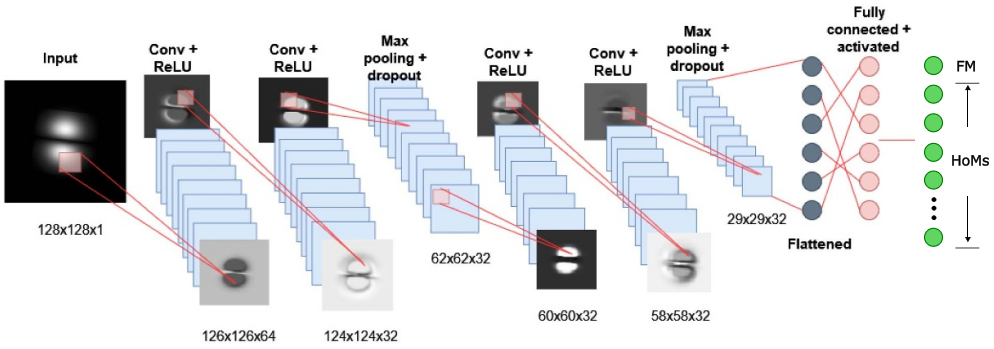

Figure 2. CNN architecture. Each layer contains a collection of feature maps of the input images. Based on the connections created between the layers during the training phase, the CNN will classify the mode images into a 0th-order (fundamental mode, i.e. FM) through to the 10th-order modes. Since the fundamental mode is chosen for determining optical alignment, we group the 1st-order through 10th-order modes together and refer to all higher-order modes as HoMs.

which determines the optimal input alignment based on features of the cavity transmission live-stream images. The result after running the algorithm is a set of input mirror positions that optimise the cavity transmission power while the cavity is resonant with the fundamental mode.

The CNN was built using Keras, a deep learning library in Python [48] to operate on a NVIDIA P100 GPU that facilitates fast parallelised computations as compared to a CPU. Figure 2 illustrates the architecture of the CNN in the system. It consists of many hidden layers, each of which performs different functions to break down the input image into its features. The CNN has 4 convolutional (conv) and rectified linear unit layers. These layers convolve the input image with various filters, allowing different image features to be identified (e.g. bright versus dark pixels). The pooling and dropout layers between the convolutional layers are implemented, which reduce the amount of data stored in the CNN [36]. After the last dropout layer, the data is flattened into a one-dimensional matrix and various activation conditions are assigned to this layer. Each condition corresponds to a 0th-order fundamental

mode (i.e. FM) through to the 10th-order mode. The 0th-order mode is chosen to determine optical alignment with the GAs. Thus, we group the 1st-order through to the 10th-order modes together and refer those higher-order modes as HoMs. In order for the CNN to effectively classify unlabelled images of the modes present in our optical cavity, it was first trained on labelled images of Hermite-Gauss cavity modes simulated using FINESSE , an interferometry simulation package in Python. The simulated mode images, generated in FINESSE, undergo various augmentations, including rotations, size adjustments, and contrast modifications, with the processing handled by Keras. Over 60 000 images of Hermite–Gaussian and Laguerre–Gaussian modes were generated and used to train and validate the CNN, 5000 of which were the fundamental mode and the rest were various higher-order modes. To obtain a more comprehensive assessment of the CNN's generalization ability beyond the validation data, a test set comprising 1200 images was utilized to evaluate the model's performance in terms of both accuracy and loss.

The CNN is combined with a GA, through which the optimal cavity input alignment is determined. The GA controls the five degrees of freedom of the optical system—the pitch (θ_x) and yaw (θ_y) angles of both the FSMs as well as the cavity length by adjusting the z position of the cavity PZT. The FSMs cover the entire range from complete misalignment, where no optical mode exists in the cavity, to optimal alignment. A conceptual diagram of our GA is shown in figure 3.

First, the algorithm generates a set of initial FSM positions, followed by the digital-to-analog converter (DAC) driving the mirrors to these positions. In the GA, the random number generation algorithm assumes a uniform distribution across the range of possible FSM positions. In the first generation, the distribution range of FSMs spans the full range of ± 26 mrad. With each subsequent generation, this range decreases by a factor of 0.6. The cavity length is scanned, using the cavity PZT, over at least one free spectral range—the frequency spacing between consecutive resonant modes of the cavity. The scan speed is set to allow the transmitted cavity modes to be captured by the camera and data acquisition system at each PZT step. The brightest images out of the scan are saved and the cavity modes are then predicted using the CNN. The optical power $P(\theta)$ is calculated from the image pixel values. If the mode is predicted to be a higher-order mode, a zero reward $R(\theta)$ is assigned to the corresponding mirror positions. If the mode is predicted as a fundamental mode, the reward $R(\theta)$ will be the power of the image captured by the camera, which is a measure of the relative power of the transmitted cavity mode. A higher fundamental mode image power is assigned with a higher $R(\theta)$, and vice versa. It was ensured that the camera was never saturated by adjusting the laser input power and a short exposure time (150 ms) for each image. After registering $R(\theta)$, the best 10% of these positions (fittest parents) are selected and mixed with 90% new random positions (the offspring) to form the next generation. In each subsequent generation, the offspring are produced though uniform crossover, a genetic operator that combines parameters from two parents in a stochastic manner. Specifically, each parameter in the offspring is selected randomly from either parent with equal probability. To enhance the quality of the offspring, the selection of higher-performing parents with higher $R(\theta)$ is prioritized. Uniform crossover treats each parameter independently, ensuring a diverse mix of information from parents. Additionally, to prevent the algorithm from becoming trapped in local maxima, random mutations are introduced by altering a single parameter in each offspring. The algorithm loop ends with optimal alignment given when a plateau in image intensity for fundamental modes is achieved in subsequent generations.

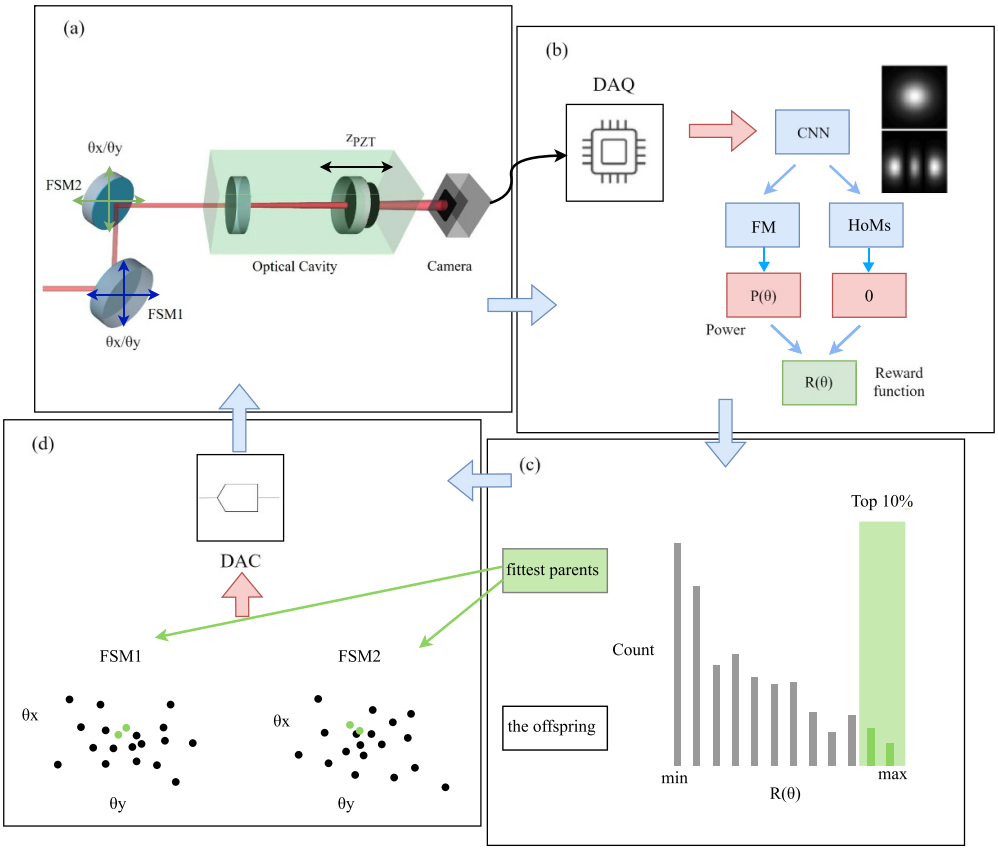

Figure 3. Conceptual GA graph. (a) Experimental setup: an incident beam is aligned to an optical cavity by two FSMs; The FSMs are controlled by an ethernet based ACROMAG XT1541 DAC module which is interfaced via our EPICS [49] control system; the DAC controls the pitch (θ_x) and yaw (θ_y) angles of the FSMs, and the cavity length z via a PZT attached. The image stream is captured by a high-speed camera and (b) passed to the data acquisition system (DAQ) for processing. The value of reward $R(\theta)$ is determined by both the CNN outcome and the image power $P(\theta)$. $R(\theta) = P(\theta)$ when the mode is fundamental mode, or $R(\theta) = 0$ when the mode is a higher-order mode. (c) The valuations of $R(\theta)$ are sorted to give and save the top 10% positions (fittest parents) in next generation (the offspring). (d) Shows a set of the selected parameters. Combined with the rest 90% being new random ones, the parameters of the next generation are fed back to experimental setup through DAC. The algorithm is continued in a loop until a termination condition satisfies.

4. Results

The CNN model in our study has a high accuracy for distinguishing real-time images of the fundamental mode and higher order modes. Two common metrics are used to evaluate the neural network in terms of its accuracy and reliability—the accuracy and loss of the model during training and validation phases, and the classification accuracy in the utilisation phase. The CNN model uses 80–20 split: the training dataset is 80% of the total data, with 20% allocated to validation. Figure 4(a) shows the accuracy and loss of the CNN model during

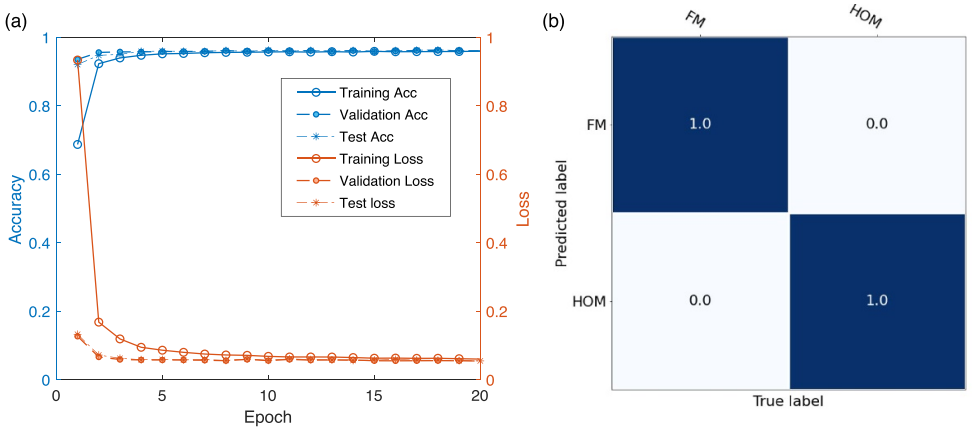

Figure 4. Convolutional neural network results. (a) CNN learning curves depicting accuracy and loss of model during training, validation and test respectively. (b) Accuracy matrix of the CNN prediction of the cavity modes (FM or HoMs). The vertical axis represents the predicted label by CNN, while the horizontal axis is the true label of the image.

the training and validation phase over 20 epochs. The CNN increases towards an accuracy of 100% by the fifth epoch, and decreases towards a loss of 0% by the tenth epoch. This indicates that the CNN learns quickly, and the model is well-fitted to the input data.

Figure 4(b) shows the classification accuracy matrix for the CNN during the utilisation phase. About 200 images of the experimental cavity transmission were provided to the CNN, 50 of which were of the fundamental mode and the rest higher order modes. The diagonals of the accuracy matrix are 1.0, which indicates that 100% of the predicted classifications are the same as the true ones. This guarantees that our alignment algorithm can converge to the best alignment for the correct mode (fundamental mode in this case).

The GA consistently determines the best beam alignment, which is also demonstrated to be reliable over multiple trials. In this study, the input beam is defined to be well-aligned when the transmitted power of the fundamental mode is maximised, meaning that the transmitted power of higher order modes is minimised. The transmitted power of the fundamental mode is quantified by the normalised intensity of the images captured by the camera to that of the input light. Starting from random mirror positions of both FSMs, the algorithm ran over 10 generations and determined the normalised intensity of the fundamental mode present in the image. With a population size of one hundred, the algorithm was always able to find fundamental modes in the first generation. By the sixth generation of the algorithm, the average intensity of the fundamental mode reaches a plateau at 90%–95% relative intensity, as shown in figure 5. Also shown in figure 5(a) are the 25th, 50th (median), and 75th percentiles of intensity to highlight both the spread and central tendency. From the sixth generation onwards, the median intensity exceeds 95%, with the 25th percentile reaching or surpassing 90%. The primary goal of optical alignment is to minimize optical losses, which is shown in figure 5(b). The algorithm reduces optical loss by a factor of 2 in more than half of the trials and achieves a reduction by one order of magnitude in 2 out of 10 trials, demonstrating a substantial enhancement in alignment quality.

As a comparison, the input beam is manually aligned to the the cavity using the fast-steering mirrors and the camera. It takes approximately $\sim$15 min for an experienced experimentalist to

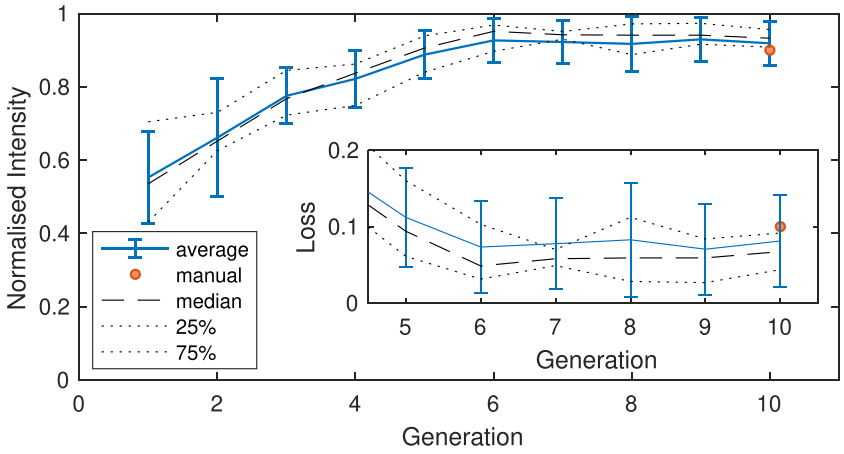

Figure 5. Normalised image intensity (a) and optical loss (b) to the input light intensity at best alignment per generation from generation 1 to 10. The image intensity is normalised to the input light intensity. The blue lines are average intensities of 10 trails. The error bars are given by the fundamental mode intensities over 10 trial runs. The dashed and dotted lines show 25th, 50th (median), and 75th percentiles of intensity to highlight both the spread and central tendency. The population for each generation is 100.

achieve a relative intensity of ∼90% manually (indicated by the red dot in figure 5). In half the time, the algorithm achieves or exceeds the alignment intensity achieved by manual methods. It is worth mentioning that algorithm evaluation prioritizes the number of generations needed and the parameter volume per generation. These metrics provide a more comprehensive assessment of the algorithm's efficiency and performance, compared to merely considering physical time.

To further qualify alignments given by the algorithm, the fluctuation of the loss in optical power due to the presence of higher order modes is estimated over the 10 trial runs. The fluctuation of the loss in optical power, $\delta P/P$ (as per figure 5), can be estimated by the couplings due to possible misalignments of an input beam with respect to the cavity fundamental modes in [14]. Table I in [14] gives coupling coefficients due to misalignment in transverse position and angular tilt as $\delta\mu/w_0$ and $\delta\theta_\mu/\theta_0$ respectively, where $\mu = x$ or y. w_0 is the cavity waist and $\theta_0 = \lambda/(\pi w_0)$ is the cavity beam divergence with the wavelength λ as 1064 nm. Thus, the fluctuation of the loss in optical power is given by $\delta P/P = 1 - (\delta x/w_0)^2 - (\delta y/w_0)^2 - (\delta\theta_x/\theta_0)^2 - (\delta\theta_y/\theta_0)^2$. Here, w_0 is 123 μm and θ_0 is 2.75 mrad. δx ($\delta\theta_x$) and δy ($\delta\theta_y$) are the standard deviations of the beam displacements (angle tilts) in the horizontal and vertical directions. The displacements (angle tilts) are calculated from propagation of the mirror tilt angles by using ABCD matrix [50], a mathematical form for performing ray tracing calculations. Over the 10 trial runs, $\delta P/P$ is calculated as 94%, meaning that there is a 6% loss of optical power from the fundamental mode (4.4% and 1.6% for optical losses caused by displacement and angular degrees of freedom respectively). This loss occurs due to fluctuations of the input beam alignment over multiple trials as well as imperfections of the mode matching between the cavity and input beam.

5. Discussion

The speed of the current experiment is largely limited by the 10 ms response time of the DAC. Over the whole algorithm, where ∼1000 mirror positions are trialled, 5 min (out of ∼15 min total) can be attributed to the DAC response time. The speed of this process can be increased by using a DAC system with a faster response time. We also employed a cost-effective GPU to demonstrate its affordability, and upgrading the GPU could also improve time efficiency.

Our demonstration system was selected for its compatibility with a wide range of optical systems, underscoring its broad applicability and versatility. The method serves as a strong foundation for developing a more systematic and comprehensive alignment solution that could be applied to a wider range of optical systems. Once initial optical alignment is achieved, angular alignment feedback and control systems can be engaged to optimise and maintain the optical cavity resonance condition. Future work will explore a more complex classification of higher-order modes, where individual modes could be analyzed and categorized independently. This finer approach may provide further improvements in performance. In future applications, the GA can be integrated with a learnable policy that would allow the algorithm to model the specific environment of the experimental setup. The statistical learning based on such a model would enable the algorithm to reach the perfect alignment much faster in the subsequent runs. Extending the scope of the algorithm to high finesse cavities, where fundamental mode may not be found within the initial generations, is another crucial direction. In such cases, the identification of higher-order modes in the transmitted beam and their utilisation in the graded reward function may be useful. The recent developments in efficient reinforcement learning is another path that can be explored for a faster agent based learning [51].

This alignment algorithm can be applied to suspended optical cavities in future work, which provides more challenges with increased number of alignment degrees of freedom and will be more susceptible to external factors such as vibrations. To enhance the robustness and generalizability of the proposed method, incorporating simulations of environmental noise and evaluating performance on larger systems would be beneficial. Reliable and optimised initial alignment of suspended optical cavities is essential for the operations of laser interferometric gravitational-wave detectors, further emphasising the need for automated alignment techniques that provide consistent but optimal results.

6. Conclusion

In this experiment, we demonstrate the use of machine learning algorithms to optimise the alignment of an optical cavity. A highly accurate CNN is utilised to classify the cavity resonant mode combined with a GA that controls the steering mirrors to optimise the input beam alignment into a fixed optical cavity. It is demonstrated that the control algorithm produces highly consistent outcomes that maximise the fundamental mode power within the cavity. Over multiple trials, the median intensity exceeds 95%, with the 25th percentile reaching or surpassing 90% from the sixth generation onwards. Overall, we show that machine learning techniques—particularly deep learning—can be utilised to automatically align an optical cavity system, producing optimal and consistent results. These results are important for high-precision optical instrumentation such as gravitational wave interferometers: machine learning can speed up the initial cavity alignment process, as well as consistently produce optical alignments to maximise the circulating power of an optical cavity.

Data availability statement

The data cannot be made publicly available upon publication because they are not available in a format that is sufficiently accessible or reusable by other researchers. The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgment

We would like to acknowledge Jack Brown and Min Jet Yap for their contributions to the framework and initial setup of this experiment.

Funding

This research is supported by the Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav), Project No. CE170100004. S J also acknowledges support from Council for Scientific and Industrial Research (CSIR), India. P C is funded by Paul Lasky and Eric Thrane via ARC LE210100002.

ORCID iDs

Jiayi Qin https://orcid.org/0000-0002-7120-9026
Shreejit Jadhav https://orcid.org/0000-0003-0554-0084
Bram J J Slagmolen https://orcid.org/0000-0002-2471-3828

References

[1] Aasi J *et al* 2015 Advanced LIGO *Class. Quantum Grav.* **32** 074001
[2] Acernese F a *et al* 2014 Advanced Virgo: a second-generation interferometric gravitational wave detector *Class. Quantum Grav.* **32** 024001
[3] Akutsu T *et al* 2019 KAGRA: 2.5 generation interferometric gravitational wave detector *Nat. Astron.* **3** 35–40
[4] Aspelmeyer M, Kippenberg T J and Marquardt F 2014 Cavity optomechanics *Rev. Mod. Phys.* **86** 1391
[5] Rosi G *et al* 2014 Precision measurement of the Newtonian gravitational constant using cold atoms *Nature* **510** 518–21
[6] Abbott B P *et al* 2019 GWTC-1: a gravitational-wave transient catalog of compact binary mergers observed by LIGO and Virgo during the first and second observing runs *Phys. Rev.* X **9** 031040
[7] Abbott R *et al* 2021 GWTC-2: compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run *Phys. Rev.* X **11** 021053
[8] Abbott R *et al* 2021 GWTC-3: compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run (arXiv:2111.03606 [gr-qc])
[9] Verhagen E *et al* 2012 Quantum-coherent coupling of a mechanical oscillator to an optical cavity mode *Nature* **482** 63–67
[10] Yap M J *et al* 2020 Generation and control of frequency-dependent squeezing via einstein–podolsky–rosen entanglement *Nat. Photon.* **14** 223–6
[11] Buikema A *et al* 2020 Sensitivity and performance of the Advanced LIGO detectors in the third observing run *Phys. Rev.* D **102** 062003
[12] Akutsu T *et al* 2020 Overview of KAGRA: detector design and construction history *Prog. Theor. Exp. Phys.* **2021** 05A101
[13] Abbott R *et al* 2020 Gw190521: a binary black hole merger with a total mass of $150M_{\odot}$ *Phys. Rev. Lett.* **125** 101102
[14] Anderson D Z 1984 Alignment of resonant optical cavities *Appl. Opt.* **23** 2944–9

[15] Du J *et al* 2013 Precision measurement of single atoms strongly coupled to the higher-order transverse modes of a high-finesse optical cavity *Appl. Phys. Lett.* **103** 083117
[16] Zhao C *et al* 2015 Parametric instability in long optical cavities and suppression by dynamic transverse mode frequency modulation *Phys. Rev.* D **91** 092001
[17] Miller J *et al* 2011 Damping parametric instabilities in future gravitational wave detectors by means of electrostatic actuators *Phys. Lett.* A **375** 788–94
[18] Capocasa E *et al* 2018 Measurement of optical losses in a high-finesse 300 m filter cavity for broadband quantum noise reduction in gravitational-wave detectors *Phys. Rev.* D **98** 022010
[19] Aasi J *et al* 2013 Enhanced sensitivity of the ligo gravitational wave detector by using squeezed states of light *Nat. Photon.* **7** 613–9
[20] Morrison E *et al* 1994 Automatic alignment of optical interferometers *Appl. Opt.* **33** 5041–9
[21] Morrison E *et al* 1994 Experimental demonstration of an automatic alignment system for optical interferometers *Appl. Opt.* **33** 5037–40
[22] Grote H *et al* 2002 The automatic alignment system of GEO 600 *Class. Quantum Grav.* **19** 1849
[23] Acernese F *et al* 2006 The Virgo automatic alignment system *Class. Quantum Grav.* **23** S91
[24] Paul J B, Lapson L and Anderson J G 2001 Ultrasensitive absorption spectroscopy with a high-finesse optical cavity and off-axis alignment *Appl. Opt.* **40** 4904–10
[25] Affeldt C *et al* 2014 Advanced techniques in GEO 600 *Class. Quantum Grav.* **31** 224002
[26] Richtmann L *et al* 2024 Model-free reinforcement learning with noisy actions for automated experimental control in optics (arXiv:2405.15421)
[27] Rakhmatulin I *et al* 2024 A review of automation of laser optics alignment with a focus on machine learning applications *Opt. Lasers Eng.* **173** 107923
[28] Sorokin D *et al* 2020 Interferobot: aligning an optical interferometer by a reinforcement learning agent *Advances in Neural Information Processing Systems* vol 33 pp 13238–48
[29] Makarenko S *et al* 2022 Aligning an optical interferometer with beam divergence control and continuous action space *Conf. on Robot Learning* (PMLR) pp 918–27
[30] Reyes-Vera E *et al* 2024 Machine learning applications in optical fiber sensing: a research agenda *Sensors* **24** 2200
[31] Gupta R K *et al* 2022 Machine learner optimization of optical nanofiber-based dipole traps *AVS Quantum Sci.* **4** 7
[32] Lei J *et al* 2022 Development and validation of a deep learning-based model for predicting burnup nuclide density *Int. J. Energy Res.* **46** 21257–65
[33] Lei J *et al* 2023 Radiation shielding optimization design research based on bare-bones particle swarm optimization algorithm *Nucl. Eng. Technol.* **55** 2215–21
[34] LeCun Y, Benigo Y and Hinton G 2015 Deep learning *Nature* **521** 436–44
[35] Sarker I 2021 Machine learning: algorithms, real-world applications and research directions *SN Comput. Sci.* **2** 160
[36] Goodfellow I, Benigo Y and Courville A 2016 *Deep Learning* (MIT Press)
[37] Ghayoumi M 2021 *Deep Learning in Practice* (CRC Press LLC)
[38] Biswas A, McIver J and Mahabal A 2020 New methods to assess and improve ligo detector duty cycle *Class. Quantum Grav.* **37** 175008
[39] Mitchell T M and Mitchell T M 1997 *Machine Learning* vol 1 (McGraw-hill)
[40] Mukund N *et al* 2023 First demonstration of neural sensing and control in a kilometer-scale gravitational wave observatory (arXiv:2301.06221)
[41] Forrest S 1993 Genetic algorithms: principles of natural selection applied to computation *Science* **261** 872–8
[42] Katoch S, Chauhan S and Kumar V 2021 A review on genetic algorithm: past, present and future *Multimedia Tools Appl.* **80** 8091–126
[43] Hluck G 1997 Genetic algorithms *The Handbook of Applied Expert Systems* ed J Liebowitz (CRC Press LLC)
[44] Sehgal A *et al* 2019 Deep reinforcement learning using genetic algorithm for parameter optimization *2019 Third IEEE Int. Conf. on Robotic Computing (IRC)* (IEEE) pp 596–601
[45] Iglesias R *et al* 2006 Combining reinforcement learning and genetic algorithms to learn behaviours in mobile robotics *ICINCO-RA* pp 188–95
[46] Chen R, Yang B, Li S and Wang S 2020 A self-learning genetic algorithm based on reinforcement learning for flexible job-shop scheduling problem *Comput. Ind. Eng.* **149** 106778
[47] Mehboob U *et al* 2016 Genetic algorithms in wireless networking: techniques, applications and issues *Soft Comput.* **20** 2467–501

[48] Ketkar N 2017 *Introduction to Keras* (Apress) pp 97–111
[49] Dalesio L R, Kozubal A and Kraimer M 1991 Epics architecture *Technical Report* Los Alamos National Lab., NM (United States)
[50] Saleh B E and Teich M C 2019 *Fundamentals of Photonics* (Wiley)
[51] Li Q, Kumar A, Kostrikov I and Levine S 2023 Efficient deep reinforcement learning requires regulating overfitting (arXiv:2304.10466)