INDIAN JOURNAL OF SCIENCE AND TECHNOLOGY

REVIEW ARTICLE

# Investigating Methods of Artificial Intelligence in Fiber Optic Network Performance Monitoring

**G Lakshmi[1]*, S Brindha[1]**

**1** Electronics and Communication Engineering, Sri Sairam Engineering College, Chennai, India

OPEN ACCESS

**Received:** 28/05/2026

**Accepted:** 28/08/2026

**Published:** 09/09/2026

**Citation:** Lakshmi G, Brindha S (2026) Investigating Methods of Artificial Intelligence in Fiber Optic Network Performance Monitoring. Indian Journal of Science and Technology 19(32): 2278-2296. https://doi.org/10.17485/IJST/v19i32.647

* **Corresponding author**.

lakshmig237@gmail.com

**Funding:** None

**Competing Interests:** None



Published By Indian Society for Education and Environment (iSee)

**ISSN**
Print: 0974-6846
Electronic: 0974-5645

## Abstract

**Background:** Elastic and cognitive optical networks vary modulation format, symbol rate, launch power and route while a service is carrying traffic, so the physical layer has to be measured continuously rather than engineered once at commissioning. Optical performance monitoring (OPM) and modulation format identification (MFI) supply that measurement, and since 2016 the preferred estimator for both has shifted from analytical models to learned ones. The surveys that mapped this shift were assembled narratively, without a published search protocol, without stated eligibility criteria, and without testing whether the accuracies they tabulated were comparable. **Objectives:** To determine which learning paradigms and architectures dominate AI-assisted OPM, MFI and fibre nonlinearity compensation and how that balance has moved; how the choice of input representation trades accuracy against computational cost; whether the reported performance figures can be pooled or ranked; which bottlenecks separate laboratory demonstrations from field deployment; and what must change in reporting practice for progress to become measurable. **Method:** Six sources were consulted, namely IEEE Xplore, Optica Publishing Group, ScienceDirect, SpringerLink, MDPI and arXiv, with Google Scholar used for backward and forward snowballing, over the window January 2014 to July 2026. Search strings combined the monitoring terms ("optical performance monitoring" OR "OSNR estimation" OR "modulation format identification" OR "nonlinearity compensation") with the learning terms ("machine learning" OR "deep learning" OR "neural network" OR CNN OR LSTM OR "support vector machine"). Records were screened against six inclusion and five exclusion criteria and appraised on a five-item, ten-point rubric covering problem definition, data provenance, train and test separation, metric definition and complexity reporting, with 6/10 as the retention threshold. Of 612 records identified, 494 were screened after de-duplication, 143 were read in full and 49 were retained. Results are presented through a PRISMA-style selection diagram, two distribution charts and ten comparative tables. **Findings:** The retained set comprises 25 primary studies and 24 reviews, tutorials and enabling works. Supervised learning accounts for 22 of the 25 primary studies (88 per cent), and no primary study applies reinforcement learning to OPM or MFI although four of the surveys propose it. By architecture the primary studies divide into feed-forward deep networks (8), convolutional networks (5), recurrent networks (5), classical and kernel methods (3), autoencoders (2), transfer learning (1) and learned multi-step compensation (1); eleven learn from image-like representations and fourteen from symbol sequences or link telemetry. Reported outcomes are strong individually and incommensurable collectively: root-mean-square errors of 0.73 dB for OSNR, 1.34 ps/nm for chromatic dispersion and 0.47 ps for differential group delay from one multi-task convolutional monitor; OSNR standard errors of 0.21 dB to 0.48 dB across four polarisation-multiplexed formats from a Stokes-space deep network; a 0.4 dB Q-factor gain from Parzen-window detection; a mean-square error of 3 x 10^-6 for an autoencoder channel model; and 42 Gb/s over 40 km from an end-to-end learned transceiver. Five different accuracy metrics are in use and no two studies share a dataset or a reference link, so these figures can be reported but not ranked. **Novelty/Significance:** This is the first review in the area to apply a documented selection protocol with explicit eligibility and quality criteria, making its corpus reproducible. It reads the literature along three axes at once, namely input representation, learning paradigm and deployability, where earlier reviews treated these separately, and it shows that representation rather than architecture governs whether a monitor is deployable. It reports a comparability audit quantifying why the published accuracies cannot be pooled, answers it with an eight-item minimum reporting set proposed here, and identifies the complete absence of reinforcement learning from the primary monitoring literature as a structural and actionable gap.

**Keywords:** Optical Performance Monitoring; Modulation Format Identification; Machine Learning; Deep Neural Networks; Fibre Nonlinearity Compensation; Elastic Optical Networks

## 1 INTRODUCTION

Optical fibre carries almost all long-haul and metropolitan traffic, and the way that traffic is carried has changed a good deal faster than the way it is measured. A decade ago a transponder was commissioned once. Its modulation format, launch power and route were fixed, and its operating margin was set by a design rule rather than by observation. Proposals for elastic and cognitive networking broke that arrangement, because format, symbol rate, spectral position and route all became variables that a control plane may change while the service is live [1,2]. Once those parameters move, the network has to know the state of its own physical layer continuously, and it has to know it without terminating the signal. Supplying that knowledge is the job of optical performance monitoring (OPM) and modulation format identification (MFI).

The classical answer to OPM was analytical. Optical signal-to-noise ratio (OSNR) was read from the noise floor between channels, chromatic dispersion and polarisation-mode dispersion were inferred from pilot tones or from delay-tap statistics, and signal quality was summarised through an error vector magnitude or a Q factor. The 2016 review by Dong and co-workers remains the best account of that generation of techniques [3]. Its own conclusions, however, anticipate the problem that followed. Out-of-band OSNR estimation loses validity as soon as a signal passes through a cascade of reconfigurable add-drop multiplexers that filter the noise along with the channel, and analytical estimators degrade further once nonlinear interference, rather than amplified spontaneous emission, dominates the noise budget. Both conditions are now normal in a dense wavelength-division multiplexed line system, which is why interest moved to estimators that learn the mapping from an observable to a parameter instead of deriving it.

Figure 1 shows the link that all of this concerns. A transmitter, a channel and a receiver, with the channel imposing the distortions that monitoring has to detect and that compensation has to undo. Everything discussed in this review is an attempt to estimate a property of that middle block, or to invert it, from observations taken at the right-hand end.

Machine learning suited the problem for a practical reason more than a theoretical one. The quantities an operator wants, namely OSNR, residual dispersion, nonlinear noise power and the format in use, are all deterministic functions of measurable signal statistics, but the functions are awkward to write down and cheap to approximate from data [4,5]. Once a receiver already performs coherent detection and digital signal processing, the observations needed to train such an approximator, whether constellation points, eye diagrams, amplitude histograms or raw symbol sequences, are available at no extra hardware cost. Reviews of network automation made the same argument at the control-plane level [6], and the trajectory from shallow classifiers to deep networks in optical communications has been documented in general terms elsewhere [7,8].

What a learned model can be asked to do depends on the shape of the data it is given, and optical links produce data of two shapes. A rendered constellation diagram, an eye diagram or an amplitude histogram is an image, and is handled by

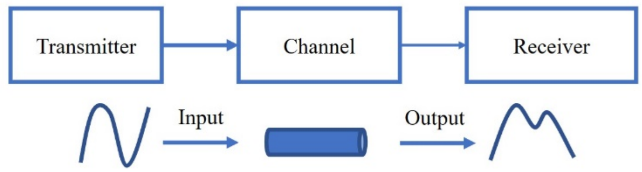


**Fig 1.** Optical Communication Link

convolutional networks. A stream of received symbols is a sequence, and is handled by recurrent networks. Figure 2 sets out the task domains that each family has been applied to in optical communication, and the division it shows runs through the whole of this review.

The exchange between the two fields runs in both directions. Neural methods are applied to photonic systems, and photonic hardware is in turn being built to carry out neural computation, which matters here because the cost of running a monitor is one of the binding constraints identified later in this review [9,10]. Figure 3 represents that mutual growth. Figure 4 places the terms used throughout in relation to one another: the architectures that dominate this corpus, namely the artificial, convolutional and recurrent neural network families, form a small part of machine learning, which is itself a small part of artificial intelligence.

The results have been substantial. A convolutional network reading constellation diagrams was shown to recover both modulation format and OSNR without any hand-designed feature extraction [11], and a deep network operating on amplitude histograms performed the same joint task in a digitally coherent receiver across QPSK, 16-QAM and 64-QAM signals [12]. A multi-task convolutional monitor reported root-mean-square errors of 0.73 dB for OSNR, 1.34 ps/nm for chromatic dispersion and 0.47 ps for differential group delay from one shared network [13]. A deep network reading density distribution in Stokes space held the standard error of OSNR estimation between 0.21 dB and 0.48 dB across four polarisation-division-multiplexed formats [14]. On the compensation side, a Parzen-window detector recovered 0.4 dB of Q factor against fibre nonlinearity [15], and an end-to-end learned transceiver reached 42 Gb/s over 40 km with transmitter and receiver optimised jointly [16]. Taken singly, each of these is a genuine advance.

What the field has not produced is an account of whether those advances add up. Three reviews dominate the citation record in this niche, and Table 1 sets out what each established and where it stops. Dong and co-workers [3] catalogued the pre-learning techniques thoroughly but predate the deep-learning wave almost entirely. Musumeci and co-workers [5] covered machine learning across the whole optical network stack, from traffic prediction to failure management, which necessarily left limited room for the monitoring layer and no room at all for comparing monitoring accuracy. Saif and co-workers [17] produced the closest predecessor to the present work, a dedicated survey of machine learning for OPM and MFI, organised by algorithm and by feature. It is a careful piece of work, but it is narrative in construction: there is no stated search protocol, no inclusion or exclusion criteria, no quality appraisal, and no examination of whether the accuracies it tabulates were obtained under conditions that make them comparable. Broader surveys of artificial intelligence in optical networking [1,2,9] share the same characteristic.

That last point is not a stylistic complaint. If a reader wants to know whether a Stokes-space deep network is a better OSNR monitor than a convolutional network reading eye diagrams, the published literature cannot answer, because the two were evaluated on different formats, at different symbol rates, over different links, and against different definitions of error. The present review therefore takes a different starting position. Rather than adding another descriptive catalogue, it selects its corpus under a documented protocol, reads that corpus along three axes simultaneously, and treats the incomparability of the reported numbers as a finding to be quantified rather than a nuisance to be smoothed over.

## 1.1 Research Contributions

The contributions of this review are of five kinds. The first is a reproducible corpus: to our knowledge this is the first review of artificial intelligence for optical performance monitoring to publish its search strings, eligibility criteria, quality-appraisal rubric and selection flow, so that a later reader can rebuild or extend the corpus rather than take it on trust. The second is a three-axis reading, in which the 25 primary studies are analysed simultaneously by input representation, by learning paradigm and by deployability; earlier reviews organise by one axis, usually algorithm, and that choice conceals the fact that representation and cost predict practicality better than architecture does. The third is a comparability audit, presented in Section 3.1.7, which demonstrates with counts that the accuracies reported in this literature involve at least five distinct metrics evaluated on non-overlapping test conditions, and states precisely what that prevents. The fourth follows from the third and is the review's own

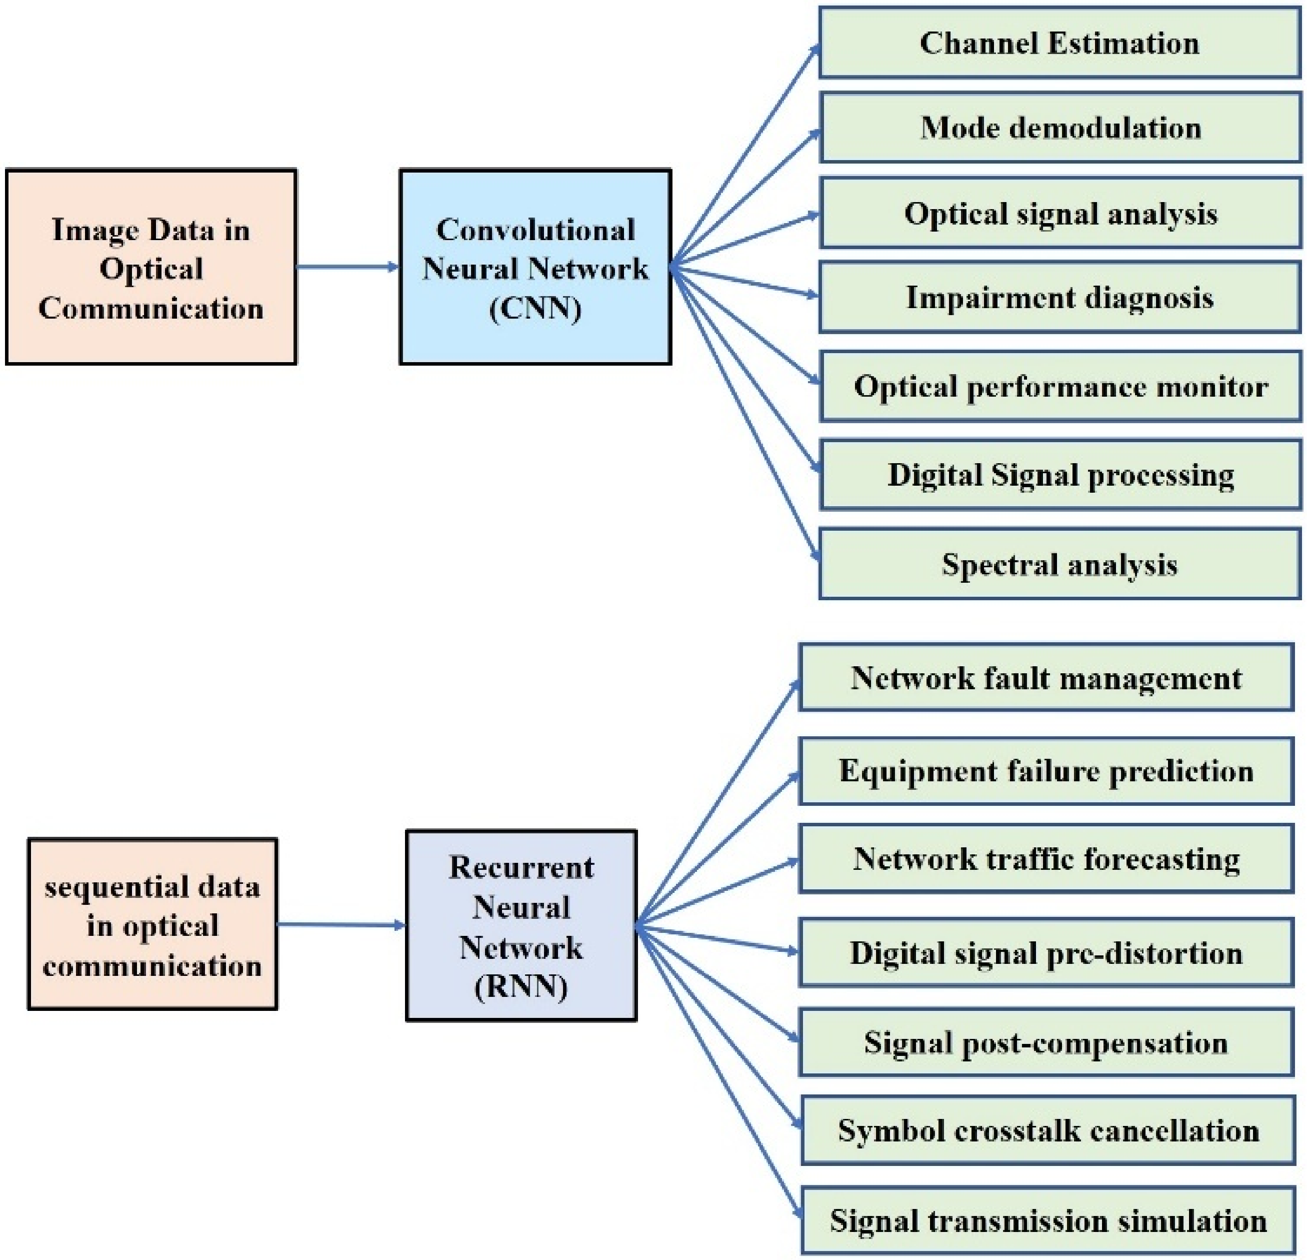


**Fig 2.** Image Data and Sequential Data Applications

proposal rather than a restatement of existing work: an eight-item minimum reporting set, items M1 to M8 in Table 10, which would make future OPM and MFI results comparable at negligible cost to the authors who adopt it. The fifth is a pair of structural observations that convert open problems into a research agenda, namely that reinforcement learning is proposed as a direction in four of the surveys read yet appears in no primary OPM or MFI study in the corpus, and that each recurring bottleneck can be traced to a root cause and matched to the line of work most likely to relieve it, as set out in Table 11.

## 1.2 Research Gaps

Reading the earlier reviews against the primary literature exposes six gaps, and they compound one another. There is no protocol: every existing review in this niche selects its material narratively, so neither its coverage nor its omissions can be

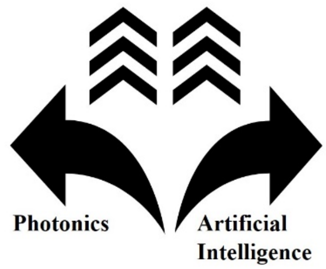


**Fig 3.** Photonics Vs Artificial Intelligence

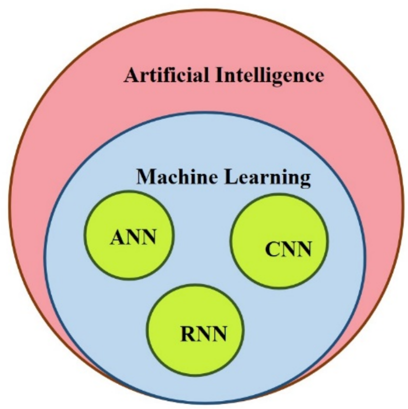


**Fig 4.** Artificial Intelligence

**Table 1.** Scope of the principal earlier reviews and the specific limitation each leaves open

| Earlier review | Scope | What it established | Limitation addressed in this review |
|---|---|---|---|
| [3] (2016) | OPM techniques for direct-detection and coherent systems | Consolidated analytical and DSP-based monitoring; identified filtering cascades and nonlinear noise as the conditions under which they fail | Predates the deep-learning literature; contains no learned estimators and therefore no basis for comparing them |
| [1] (2018) | Artificial intelligence across optical networking | Mapped AI use cases from the physical layer to network management | Breadth over depth: OPM and MFI appear as one-use case among many, with no quantitative comparison |
| [5] (2019) | Machine learning across the optical network stack | Established the taxonomy of ML tasks in optical networking and the data each requires | Monitoring accuracy is described but not compared; no eligibility criteria or appraisal of the studies cited |
| [4] (2019) | Tutorial perspective on ML for optical communications | Explained why the OPM problem is well matched to supervised regression and classification | Tutorial in intent; selection of illustrative work is not systematic |
| [17] (2020) | Dedicated survey of ML for OPM and MFI | The most complete catalogue of techniques and features in this niche to date | Narrative selection with no protocol, criteria or quality appraisal; tabulated accuracies are not tested for comparability; complexity is discussed qualitatively |
| [2] (2020) | AI-based modelling and monitoring for elastic optical networks | Linked monitoring to elastic network operation and control | Forward-looking rather than retrospective; no systematic corpus |

checked. There is no comparability: accuracies are tabulated side by side as though they were commensurable, when the underlying test conditions differ in format, symbol rate, link length and noise loading, and the error definitions differ as well. Cost is described rather than measured, although training time, inference latency, memory bandwidth and arithmetic complexity are exactly the properties that decide whether a monitor can sit inside a live receiver; the memory-bandwidth limit reported for recurrent monitors [18] and the branch-count growth reported for deep equalisers [19] are the exceptions that prove the point. Simulation and experiment are conflated, with results from split-step simulation environments listed alongside results measured on laboratory hardware even though the generalisation risk attached to each is quite different. Training data is treated as free, since nearly every study assumes a labelled dataset spanning the operating range, when acquiring one from a production network is the single largest practical obstacle and only a small part of the literature engages with it [20]. Finally, decision-making is absent: the monitoring literature stops at estimation, what an operator does with an estimate is a sequential decision problem left to the control-plane literature, and the two have not been joined.

### 1.3 Research Questions

Those gaps translate into five questions, which the remainder of the review answers in order.

RQ1. Which learning paradigms and network architectures dominate AI-assisted OPM, MFI and nonlinearity compensation, and how has the balance shifted across the period covered? (Section 3.1.1, Section 3.1.2)

RQ2. How does the choice of input representation, image-like against sequence-like, affect achievable accuracy and computational cost? (Section 3.1.3)

RQ3. Are the performance figures reported in this literature comparable, and if not, what specifically prevents comparison? (Section 3.1.4, Section 3.1.7)

RQ4. Which bottlenecks stand between the published results and a deployable monitor, and how has the literature addressed them? (Section 3.1.5, Section 3.1.6, Section 3.1.8)

RQ5. Which directions are most likely to yield field-deployable OPM and MFI, and what must change in reporting practice for progress towards them to be measurable? (Section 3.1.7, Section 3.1.8, Section 4)

### 1.4 Organisation of the Review

Section 2 states the review methodology: the sources searched, the strings used, the eligibility criteria, the quality rubric and the selection flow. Section 3 opens with an overview of how the synthesis proceeds and then presents the findings in eight subsections, moving from the composition of the corpus, through the milestone results and the representation and cost trade-off, to the comparability audit, the proposed reporting set and the bottleneck map. Section 4 draws the conclusions, states the limitations of the review itself and sets out the knowledge gaps that remain. Table 2 lists the acronyms used throughout.

**Table 2.** Acronyms used in this review

| Acronym | Expansion | Acronym | Expansion |
|---|---|---|---|
| AI | Artificial Intelligence | NLC | Fibre Nonlinearity Compensation |
| ANN | Artificial Neural Network | OPM | Optical Performance Monitoring |
| BER | Bit Error Rate | OSNR | Optical Signal-to-Noise Ratio |
| CD | Chromatic Dispersion | PAM | Pulse Amplitude Modulation |
| CNN | Convolutional Neural Network | PDM | Polarisation-Division Multiplexing |
| DGD | Differential Group Delay | QAM | Quadrature Amplitude Modulation |
| DNN | Deep Neural Network | QoT | Quality of Transmission |
| DSP | Digital Signal Processing | QPSK | Quadrature Phase-Shift Keying |
| EVM | Error Vector Magnitude | RMSE | Root Mean Square Error |
| FSO | Free-Space Optics | RNN | Recurrent Neural Network |
| LSTM | Long Short-Term Memory | ROADM | Reconfigurable Optical Add-Drop Multiplexer |
| MFI | Modulation Format Identification | SVM | Support Vector Machine |
| MFR | Modulation Format Recognition | VLC | Visible Light Communication |
| ML | Machine Learning | WDM | Wavelength Division Multiplexing |

## 2 REVIEW METHODOLOGY

The corpus was assembled under a written protocol, fixed before screening began, so that the selection can be audited and repeated. The protocol follows the structure of a PRISMA workflow, adapted to an engineering literature in which conference proceedings and preprints carry real weight and in which formal effect sizes do not exist.

### 2.1 Sources and Search Strings

Six bibliographic sources were searched, chosen because between them they index essentially all venues in which optical communication and machine learning intersect. Google Scholar was not used as a primary source, because its coverage is not reproducible, but it was used for backward and forward snowballing from the records already retained. Every search combined one monitoring term with one learning term, using the Boolean form (”optical performance monitoring” OR ”OSNR estimation” OR ”OSNR monitoring” OR ”modulation format identification” OR ”modulation format recognition” OR ”nonlinearity compensation” OR ”nonlinear equalization”) AND (”machine learning” OR ”deep learning” OR ”neural network” OR ”convolutional neural network” OR ”recurrent neural network” OR ”support vector machine” OR autoencoder). Searches were run against title, abstract and keyword fields, restricted to the window January 2014 to July 2026. Table 3 gives the yield by source.

**Table 3.** Sources searched and records retrieved

| Source | Principal coverage relevant to this review | Records retrieved |
|---|---|---|
| IEEE Xplore | JLT, Photonics Technology Letters, Photonics Journal, Access, Communications Surveys and Tutorials, OFC/ECOC/ICTON/ICOCN proceedings | 268 |
| Optica Publishing Group | Optics Express, JOCN, OFC digest | 96 |
| ScienceDirect (Elsevier) | Optical Switching and Networking, Optics Communications | 74 |
| SpringerLink | Photonic Network Communications, Optical and Quantum Electronics | 58 |
| MDPI | Applied Sciences, Photonics, Sensors | 41 |
| arXiv | eess.SP and cs.LG preprints on optical channel learning | 45 |
| Snowballing (Google Scholar) | Backward and forward citation tracing from retained records | 30 |
| **Total identified** | | **612** |

### 2.2 Eligibility Criteria

Criteria were applied in two passes, first to titles and abstracts and then to full texts. They are stated in Table 4. The inclusion rules are deliberately narrow on one point: a record was kept only if it reported a quantitative outcome, whether an error, an accuracy, a gain or a complexity figure. Position papers and purely conceptual pieces were excluded from the primary set, although a small number were retained as background where they define the terms the primary studies use.

### 2.3 Study Selection

Figure 5 shows the flow. Of 612 records identified, 118 were duplicates, leaving 494 for title and abstract screening. Screening removed 351, most of them records in which the phrase ”neural network” referred to photonic hardware for computing rather than to a learned estimator of a transmission parameter. The 143 remaining records were read in full, and 94 were excluded with the reasons recorded in the figure. Forty-nine records were retained.

### 2.4 Quality Appraisal

Each full-text record was scored on five items, 0 for absent, 1 for partial and 2 for complete, giving a maximum of 10. The rubric is given in Table 5. Records scoring below 6 were excluded, which removed 17 of the 143 full-text records. The 49 retained records had a mean score of 7.4 out of 10. The distribution of that score is itself informative. Items Q1 to Q4 were satisfied fully

**Table 4.** Inclusion and exclusion criteria

| Inclusion criteria | Exclusion criteria |
|---|---|
| I1. Addresses OPM, MFI, quality-of-transmission estimation, fibre nonlinearity compensation or optical-layer failure detection. | E1. No learning component; the method is analytical or conventional DSP only. |
| I2. Applies a learning-based method, whether supervised, unsupervised, self-supervised or reinforcement learning. | E2. Application domain is outside optical transmission and its close neighbours (visible light and free-space optics were retained; radio-only work was not). |
| I3. Reports at least one quantitative outcome (error, accuracy, gain, complexity or latency). | E3. No quantitative outcome reported. |
| I4. Published between January 2014 and July 2026. | E4. Superseded by an extended version from the same group, in which case only the extended version was kept. |
| I5. Full text available in English. | E5. Quality-appraisal score below 6 out of 10 (Section 2.4). |
| I6. Peer-reviewed journal or conference paper, or a preprint that is clearly the primary record of the work. | |

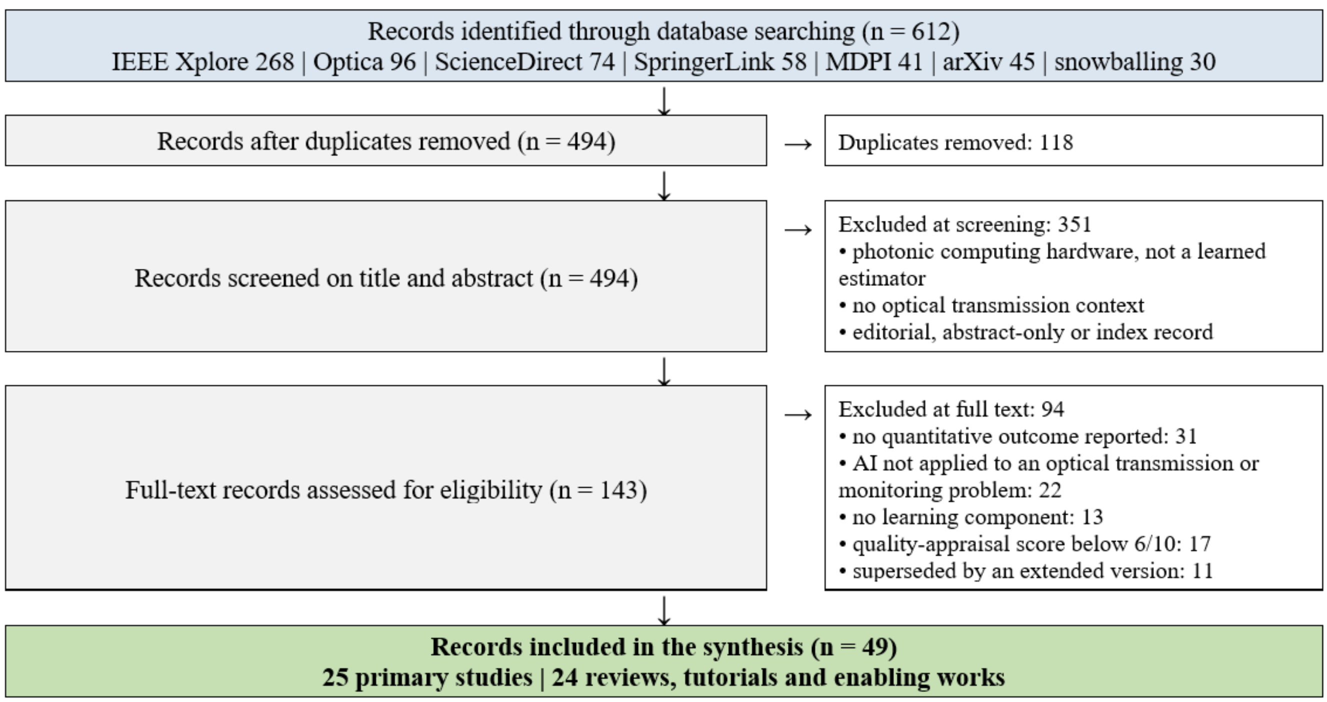


**Fig 5.** PRISMA-style selection flow for the review corpus

by most retained records, whereas Q5, which concerns computational cost, was the item on which the corpus scored worst: complexity is far more often asserted than measured. That observation is developed in Section 3.1.8.

## 2.5 Data Extraction and Synthesis

For every retained record we extracted the monitoring target, the algorithm family, the input representation, the evaluation environment, the reported quantitative outcome and the limitation stated by the authors themselves. Records were then classified as primary studies, meaning work that reports an original learned estimator with a measured outcome, or as reviews, tutorials and enabling works, meaning surveys, dataset and infrastructure papers, and studies in adjacent domains that the primary literature depends on. Synthesis is narrative and comparative rather than statistical. A meta-analysis in the formal sense is not possible here, and Section 3.1.7 explains exactly why, which is one of the results of this review rather than an apology for its method.

**Table 5.** Quality-appraisal rubric applied to full-text records

| Item | Criterion | Scoring |
|---|---|---|
| Q1 | The monitoring target and the operating range are stated unambiguously | 2 = target and range given; 1 = target only; 0 = neither |
| Q2 | Data provenance is described: simulation platform or experimental testbed, symbol rate, format, link length, channel count | 2 = all reported; 1 = partially; 0 = not described |
| Q3 | Training, validation and test data are separated, and the separation is stated | 2 = explicit separation; 1 = implied; 0 = not addressed |
| Q4 | Performance is reported against a defined metric, with the operating range over which it holds | 2 = metric and range defined; 1 = metric only; 0 = qualitative claim |
| Q5 | Computational cost is addressed: parameter count, arithmetic operations, training time, inference latency or memory requirement | 2 = quantified; 1 = discussed qualitatively; 0 = not addressed |

# 3 OVERVIEW

This section presents the outcome of the synthesis. It proceeds in a fixed order, and it is worth stating that order in advance because each step depends on the one before it. Section 3.1.1 characterises the corpus itself, by year, by paradigm and by architecture, which answers the first half of RQ1 and establishes the population from which every later claim is drawn. Section 3.1.2 identifies the milestone contributions and explains what each one settled, which answers the second half of RQ1. Section 3.1.3 turns to input representation, the axis that in our reading separates practical monitors from impressive ones, and answers RQ2. Sections 3.1.4 to 3.1.6 present the quantitative results by task, respectively OSNR and format monitoring, nonlinearity compensation, and failure detection and network automation. Section 3.1.7 then audits those results for comparability, answers RQ3, and proposes the reporting set that follows from the audit. Section 3.1.8 consolidates the bottlenecks and maps them to emerging directions, answering RQ4 and RQ5. Throughout, a claim about the corpus is expressed as a count out of a stated denominator, so that a reader can check it against the reference list.

## 3.1 Results and Findings

### *3.1.1. Composition of the corpus*

The 49 retained records divide into 25 primary studies and 24 reviews, tutorials and enabling works. Figure 6 gives the year-wise distribution of the full set. Two features stand out. Activity roughly doubles between 2017 and 2018, which coincides with the first constellation-image and amplitude-histogram monitors reaching print [11,12,21], and it peaks in 2020, the year in which the dedicated survey of the field appeared [17] and in which the largest number of primary results were published. The distribution is not a claim about the field as a whole; it is the distribution of the records that passed the eligibility and quality criteria in Section 2.


**Fig 6.** Year-wise distribution of the 49 retained records

Figure 7 breaks the 25 primary studies down by architecture. Feed-forward deep networks are the largest group with 8 studies, followed by convolutional networks and recurrent networks with 5 each, then classical and kernel methods with 3, autoencoders with 2, and single instances of transfer learning and of learned multi-step compensation. The paradigm split is more lopsided than the architecture split: 22 of the 25 primary studies (88 per cent) train under supervision, 2 use autoencoders in an unsupervised or self-supervised arrangement [22,23], and 1 learns compensation parameters by gradient descent without a classification or regression target [24].

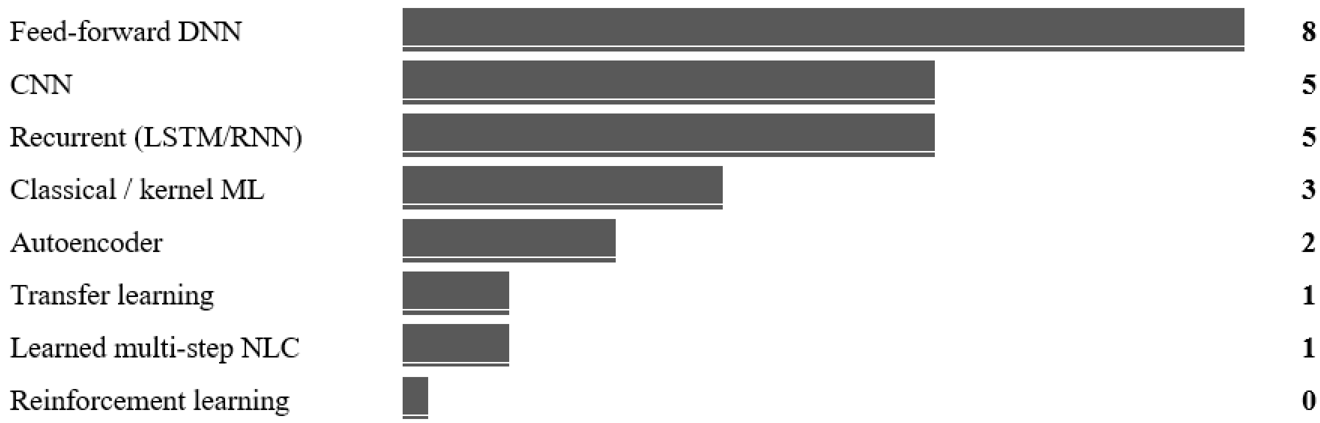


**Fig 7.** Architecture distribution across the 25 primary studies

The last row of Figure 7 is the one worth pausing on. Reinforcement learning is put forward as a direction for optical networking in four of the surveys we read [1,4–6], and it is the natural formalism for the problem that monitoring exists to serve, namely deciding what to do when a parameter drifts. Yet no primary study in the retained corpus applies it to OPM or MFI. The reason is structural rather than accidental. Reinforcement learning needs an environment that can be acted upon and that returns a reward, whereas a monitor is a passive estimator installed on a link that an operator will not allow an untrained agent to perturb. Until a validated digital twin of a line system exists, and the channel-modelling work [25] and automated data-collection work [20] in the corpus are the closest approaches to one, the gap will persist. We return to this in Section 3.1.8.

For orientation, Figure 8 shows the standard division of learning paradigms as it applies here. Supervised learning maps an input to a known output through a function learned from labelled examples, which in this literature means mapping a signal representation to an OSNR value or a format label. Unsupervised learning finds structure without labels and appears in the corpus mainly as autoencoder-based representation learning [23,26]. Reinforcement learning learns a policy from the consequences of actions, and is, as noted, absent.

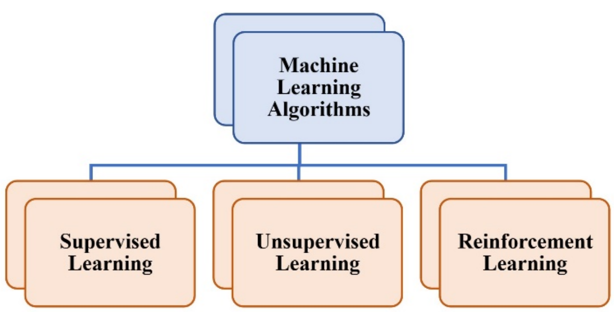


**Fig 8.** Types of Machine Learning Algorithms

Figure 9 opens the first two of those categories out into the individual algorithm families. Set against the counts in Figure 7, it makes the imbalance in this corpus plain: of the seven supervised families listed, only four appear in the primary studies at all, the convolutional, recurrent and feed-forward networks together with the kernel methods, while decision trees, random forests and nearest-neighbour classifiers survive almost exclusively as baselines that a neural model is reported to beat [21]. On the unsupervised side, clustering does not appear in a single primary study, and dimensionality reduction appears only in the autoencoder work [22,23].

It is worth being clear about what does and does not differ between the architectures counted in Figure 7. The computational element underneath all of them is the same weighted sum followed by a nonlinearity, shown in Figure 10. What separates a

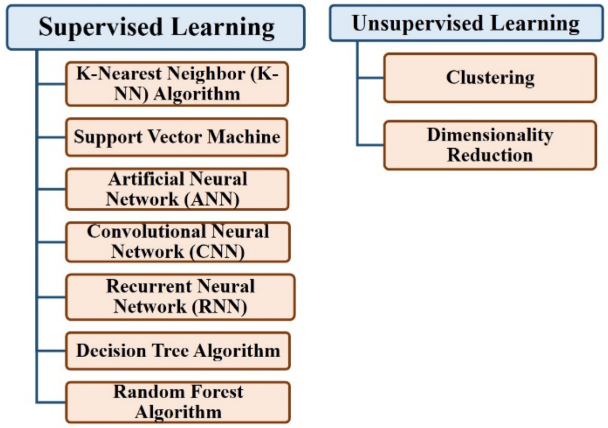


**Fig 9.** Algorithms of Supervised and Unsupervised Learning

convolutional monitor from a recurrent one is not that element but how the elements are wired and what the network is asked to consume, and it is that second choice, the input representation, which Section 3.1.3 identifies as the property that actually governs deployability.

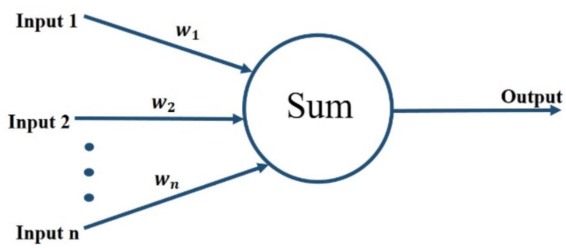


**Fig 10.** Single Neuron Structure

### *3.1.2. Milestone contributions*

Not every paper in a corpus of this kind changes what follows it. Table 6 identifies the eight contributions that did, in the sense that later work in the corpus builds on them, adopts their representation or replicates their experiment. Grouping them chronologically shows the field moving through three phases. Between 2016 and 2017 the question was whether a learned estimator could match an analytical one at all, and the constellation-image and amplitude-histogram monitors answered it. Between 2018 and 2019 the question became whether one model could serve several monitoring targets at once, which multi-task learning[13] and joint format-and-OSNR estimation[12,14] settled affirmatively. From 2020 the question shifted again, to whether these estimators can be made cheap and general enough to deploy, and that question is still open.

### *3.1.3. Input representation and its cost*

The axis along which this literature divides most cleanly is not the algorithm but the representation fed to it. Eleven of the 25 primary studies learn from an image-like representation of the signal, and fourteen learn from a symbol sequence or from link telemetry. The distinction has consequences that the algorithm choice does not.

Image-like representations are attractive because they let a convolutional network do the feature engineering. A constellation diagram, an eye diagram, an amplitude histogram, a Stokes-space density map or a high-resolution optical spectrum can all be rendered as a two-dimensional array, and the network then learns which regions of that array carry information about the parameter of interest. The intelligent constellation analyser[11] is the clearest demonstration: it takes raw pixel values, needs no prior knowledge of the eye-diagram parameters or of the transmitted bits, and still recovers format and OSNR, alongside error vector magnitude and bit-error-rate estimation. Convolutional monitors reading eye diagrams reached comparable accuracy

**Table 6.** Milestone contributions in the corpus and what each settled

| Year | Work | Contribution | Why it mattered for what followed |
|---|---|---|---|
| 2016 | [3] | Consolidated the pre-learning OPM techniques and their failure conditions | Defined the problem the learning literature then attacked; still the reference point for what a learned monitor has to beat |
| 2017 | [11] | Convolutional network reading constellation diagrams as images, recovering format and OSNR jointly | Established the image-of-the-signal representation that a third of the later primary studies adopt |
| 2017 | [12] | Two-stage deep network: format identification from amplitude histograms, then a format-specific OSNR monitor | Showed that staging beats a single monolithic model, and that no transmitter or hardware change is required |
| 2018 | [16] | End-to-end deep learning of the complete transceiver, 42 Gb/s over 40 km | Opened the line of work that treats the whole link, rather than one receiver block, as the object to be learned |
| 2018 | [13] | Multi-task convolutional monitor: OSNR, chromatic dispersion, differential group delay and bit-rate/format from one network | Demonstrated that shared representations reduce the per-parameter cost of monitoring, the strongest argument yet for learned OPM |
| 2019 | [14] | Stokes-space density distributions as the input representation for joint MFI and OSNR monitoring | Introduced a representation that is compact and modulation-agnostic, avoiding the pixel cost of image methods |
| 2020 | [24] | Machine-learning optimisation of multi-step nonlinearity compensation, demonstrated experimentally at 25 GBaud | Reframed complexity reduction as a design problem rather than a step-count problem |
| 2020 | [17] | Dedicated survey of ML for OPM and MFI | Consolidated the field and, by its own construction, exposed the need for the protocol-driven treatment attempted here |
| 2024 | [20] | Comprehensive review of ML/DL applications in short-reach optical systems, including OPM, MFI, equalization | Established a framework for comparing AI methods and highlighted key challenges—data scarcity, computational complexity, generalization |

and outperformed decision trees and support vector machines on the same inputs[21], and the same approach extends to grayscale constellation images across eight digital formats when the network is deepened with cascaded flow-in-flow blocks[27].

The cost is paid at inference (Table 7). Rendering a constellation or eye diagram means accumulating enough symbols to populate the image, then convolving over a two-dimensional array whose pixel count bears no relation to the information it carries. Every one of the image-based studies in the corpus that discusses timing at all notes that feature extraction from optical images takes longer than the algorithm itself. That is the reverse of the usual assumption about learned monitors, and it is the reason the compact representations matter: Stokes-space densities[14] and optical spectra[28] both carry the OSNR-relevant information in far fewer numbers than a rendered image does. The amplitude-histogram approach[12] sits between the two, cheap to form but format-dependent, which is precisely why that work needed a format-identification stage before the monitoring stage.

Sequence-like representations avoid the rendering cost and buy a different problem. Recurrent models operating on symbol sequences can exploit correlations between past and present samples, which is exactly what fibre memory produces, and they are therefore the natural fit for nonlinear noise estimation[18] and for post-equalisation[19]. The reported obstacle is memory bandwidth rather than arithmetic: the linear layers in a long short-term memory network demand enough bandwidth that compute units sit idle waiting for weights[18]. This is a hardware-level constraint, and it does not improve with a better training procedure. The split between the two families, and the task domains each has been applied to, is the one shown in Figure 2.

### *3.1.4. OSNR monitoring and modulation format identification*

This is the densest part of the corpus and the part where the comparability problem is sharpest. Table 8 collects the primary OSNR and MFI studies with the outcome each reports, expressed in the units the authors used.

Read as a set, three things are apparent. First, the strongest results come from models that solve several monitoring problems at once rather than one. The multi-task convolutional monitor[13] returns OSNR, chromatic dispersion and differential group delay from a single network at 0.73 dB, 1.34 ps/nm and 0.47 ps root-mean-square error respectively, and does so while also identifying bit rate and format. Treating monitoring as one estimation problem with several outputs is more accurate and cheaper than treating it as several problems, and this is the most transferable finding in the corpus.

Second, staging works. Khan and co-workers[12] identify the format first with one deep network operating on amplitude histograms, then select a format-specific OSNR monitor from a bank of networks trained per format. This is more accurate than asking one model to be format-agnostic, and it costs nothing in hardware because both stages run on data the coherent

**Table 7.** Input representations, their properties and their cost

| Representation | Used by | Strength | Cost or limitation |
|---|---|---|---|
| Constellation diagram (image) | [11,23,27] | No hand-designed features; carries format and noise information simultaneously | Needs enough symbols to populate the image; pixel count unrelated to information content; rendering dominates inference time |
| Eye diagram | [18,21] | Directly interpretable; well matched to convolutional feature extraction | Same rendering cost; error-free OSNR estimation reported only at high epoch counts |
| Amplitude histogram | [12] | Very cheap to form; no phase recovery needed | Format-dependent, so a format-identification stage must precede monitoring |
| Stokes-space density | [14] | Compact and modulation-agnostic; strong separation between PDM formats | Requires polarisation-resolved coherent reception |
| Optical spectrum | [28,29] | In-band OSNR without demodulation; works with kernel methods on small datasets | Needs high-resolution spectral measurement; degrades under tight filtering |
| Phase portrait | [30] | Preserves temporal structure in a two-dimensional form; suits transfer learning across formats | Interpretation is less direct than a constellation diagram |
| Symbol sequence (time domain) | [18,19,31–33] | Exploits channel memory; no rendering step | Memory-bandwidth bound in recurrent layers; complexity grows with the number of taps or branches |
| Link telemetry and power spectra | [34,35] | Available from existing network management systems | Data collection is slow; labels depend on rare failure events |

receiver already produces. The same principle reappears in cascaded form with transfer learning between stages[30], which reduces the retraining burden when a new format is added.

Third, and least comfortably, the numbers in Table 8 cannot be ranked. A root-mean-square error in decibels, a standard error of the mean in decibels, a classification accuracy across eight formats and a mean-square error on a channel model are four different quantities. Even where two studies report the same quantity, they report it over different formats, different symbol rates and different link lengths. This is developed in Section 3.1.7.

### *3.1.5. Nonlinearity compensation and equalisation*

Fibre nonlinearity sets the capacity ceiling of a wavelength-division-multiplexed system, and it is deterministic, which makes it a natural target for a learned inverse. The corpus contains ten primary studies here, condensed into six lines of Table 9 because the end-to-end learning line is treated as one group, and they fall into two camps that are rarely distinguished in earlier reviews.

The first camp treats the learned model as a detector or equaliser placed after conventional digital signal processing. Amari and co-workers[15] applied a Parzen-window classifier as a detector in both dispersion-managed and unmanaged systems and recovered 0.4 dB of Q factor. Zhao and co-workers[31] used a simple recurrent network with explicit time memory on single-channel and three-channel WDM PDM-16QAM systems at 62.5 GHz spacing, arguing that the historical information a recurrent unit retains is what the nonlinear channel actually requires. Deligiannidis and co-workers[32] evaluated recurrent variants specifically for complexity as well as accuracy, which is the exception in this corpus rather than the rule. In visible-light systems the same idea appears as a memory-controlled deep LSTM post-equaliser[19], where the practical finding is blunt: computational complexity grows with the number of branches, so the equaliser has to restrict short-term memory to the final taps to stay within budget.

The second camp treats learning as a way of optimising a physically motivated compensator rather than replacing it. Oliari and co-workers[24] make the argument most clearly. Their experimental demonstration on a single-channel 25 GBaud system shows that complexity reduction need not come from cutting the number of digital back-propagation steps; it can come from designing and optimising the steps better, or from decomposing expensive operations into cheaper ones without losing performance. This is a more conservative use of machine learning and, on the evidence of the corpus, a more transferable one, because the resulting compensator retains the structure of a physical model and therefore generalises beyond the conditions it was trained on. The end-to-end learning line[16,36–38] pushes in the opposite direction, optimising transmitter and receiver jointly, and the convolutional-recurrent hybrid[33] sits between the two. Table 9 summarises the group.

**Table 8.** Primary OSNR monitoring and modulation format identification studies

| Work | Target | Algorithm | Input | Reported outcome | Stated limitation |
|---|---|---|---|---|---|
| [11] | MFR and OSNR | CNN | Constellation diagram | Joint format recognition and OSNR estimation from raw pixels; also supports EVM, BER and linear/nonlinear fault monitoring | Rendering and feature extraction dominate processing time |
| [21] | OSNR | CNN | Eye diagram | Higher accuracy than decision tree and SVM on identical inputs | Error-free estimation only at high epoch counts |
| [12] | OSNR and MFI | DNN (two-stage) | Amplitude histogram | Joint monitoring for QPSK, 16-QAM and 64-QAM in a digital coherent receiver; no transmitter or hardware modification | Training consumes substantial computing resources and time |
| [18] | OSNR and nonlinear noise power | LSTM | Constellation and eye diagram | Simultaneous OSNR and nonlinear noise estimation on a 5-channel long-haul system, -3.0 to +3.0 dBm per channel (VPI simulation) | Linear layers are memory-bandwidth bound; compute units under-used |
| [13] | OSNR, CD, DGD, bit rate and format | CNN with multi-task learning | Image-like signal representation | RMSE of 0.73 dB (OSNR), 1.34 ps/nm (CD) and 0.47 ps (DGD) from one shared network | Joint optimisation couples the tasks; a change in one target requires retraining all |
| [14] | MFI and OSNR | DNN | Stokes-space density distribution | Standard error of the estimated mean OSNR between 0.21 dB and 0.48 dB across PDM-QPSK, PDM-8QAM, PDM-16QAM and PDM-32QAM | Requires polarisation-resolved coherent reception |
| [28] | In-band OSNR | SVM and Gaussian process | High-resolution optical spectrum | In-band estimation demonstrated on a 28 GBd PM-QPSK signal generated with a tunable laser at 1550.9 nm | Depends on spectral resolution; sensitive to filtering cascades |
| [30] | MFI and OSNR | Cascaded NN with transfer learning | Phase portrait | Joint identification and monitoring with reduced retraining when formats are added | Transfer performance depends on similarity between source and target formats |
| [27] | Modulation classification | Deep CNN (flow-in-flow blocks) | Grayscale constellation image | Classification across eight digital modulation formats on a deliberately difficult dataset | Depth increases inference cost; evaluated on synthetic imagery |
| [29] | Multi-parameter spectral analysis | ML on spectral features | Optical spectrum | Single spectral front end serving several monitoring functions | Accuracy tied to spectral resolution |

Two adjacent lines support this group without belonging to it. Learned channel models [25] and autoencoder-based constellation shaping [23] both aim at the same object, a differentiable surrogate for the fibre, and the sparse-autoencoder channel estimator for multi-wavelength visible-light systems reports a mean-square error of 3 x 10^-6, the lowest of the alternatives its authors compared [22]. Work on optical hardware for neural computation [10,39,40], on the use of deep learning to design the photonic devices themselves [41], and on deep detection for free-space links [42] sits further out still, but explains why the field expects the inference cost problem to become tractable.

### *3.1.6. Failure detection and network-level automation*

Monitoring exists so that something can be done with the result, and a small but practically important part of the corpus closes that loop. Lun and co-workers [35] proposed a two-stage soft-failure detection scheme combining a one-dimensional convolutional network with the receiver digital signal processing, addressing the class of gradual degradations that do not break a link but quietly erode its margin. Because soft failures are becoming more common as systems become more dynamic, and because early identification is what allows a degraded link to be restored before service is interrupted, this is arguably the most operationally relevant task in the corpus.

**Table 9.** Nonlinearity compensation and equalisation studies

| Work | Approach | System evaluated | Reported outcome and limitation |
|---|---|---|---|
| [15] | Parzen-window classifier as detector | Dispersion-managed and unmanaged links | 0.4 dB Q-factor improvement. Acts as a detector rather than a channel inverse, so gains do not compound with DSP-based compensation |
| [24] | ML-optimised multi-step compensation | Single-channel 25 GBaud, experimental | Complexity reduced by better step design and operation decomposition rather than by fewer steps. Retains physical structure, which aids generalisation |
| [31] | Simple recurrent network with time memory | Single-channel and 3-channel WDM PDM-16QAM, 62.5 GHz spacing (VPI simulation) | Presented as a low-complexity alternative to ANN and DNN compensators. Validated in simulation only |
| [32] | Recurrent network variants | Digital coherent systems | Evaluates complexity alongside accuracy, one of the few studies in the corpus to do so |
| [19] | Memory-controlled deep LSTM post-equaliser | High-speed PAM visible-light system | Compensates linear and nonlinear distortion and inter-symbol interference. Complexity grows with branch count, forcing short-term memory to be restricted to the final taps |
| [33] | Convolutional-recurrent hybrid | Nonlinear fibre-optic link | Combines local feature extraction with sequence memory to reduce parameter count |
| [16,36–38] | End-to-end learned transceiver | Intensity-modulated and long-haul coherent links | 42 Gb/s over 40 km in the original demonstration. Requires a differentiable channel model, which limits transfer to deployed links |

At the network layer, Mo and co-workers[34] trained a deep network on data collected from a 90-channel reconfigurable add-drop multiplexer system to recommend wavelength assignments that minimise power excursions during switching. The result is useful and the paper is honest about the cost: collecting the training data was slow. That admission is the important one, because it generalises. Related work on optical-layer recovery for critical services[43], on security management[44], on quality-of-transmission estimation[45,46] and on the practical implementation of monitoring techniques[47,48] shares the same dependency: the model is straightforward, and the labelled data is not. The automated dataset-collection system of Lu and co-workers[20] is the only work in the corpus that treats data acquisition as the primary research problem, and it deserves more attention than it has received.

### *3.1.7. Comparability audit and a proposed minimum reporting set*

This subsection answers RQ3 directly, and it is where this review departs most sharply from its predecessors. Earlier surveys present tables of accuracies from different studies in adjacent rows, which invites the reader to compare them. We tested whether that comparison is legitimate. It is not, for four separate reasons, each of which we can state as a count.

First, metric heterogeneity. Across the primary studies the reported accuracy appears in at least five incompatible forms: root-mean-square error in decibels and in dispersion units[13], standard error of the estimated mean in decibels[14], Q-factor gain in decibels[15], mean-square error on a modelled quantity[22], and classification accuracy over a format set[27]. A root-mean-square error and a standard error of the mean are not the same statistic and do not convert into one another without the underlying distribution.

Second, test-condition heterogeneity. No two primary studies in the corpus share a link. The conditions include a 5-channel long-haul system swept from -3.0 to +3.0 dBm per channel[18], a three-channel WDM system at 62.5 GHz spacing[31], a single-channel 25 GBaud experimental link[24], a 28 GBd PM-QPSK signal at 1550.9 nm[28] and a 90-channel ROADM testbed[34]. An OSNR estimator that is accurate over one of these has not been shown to be accurate over another.

Third, format-set heterogeneity. The format sets range from three[12] to four[14] to eight[27]. Classification accuracy over three formats and over eight formats are not comparable numbers, and the difficulty of the task scales with the set.

Fourth, the simulation and experiment distinction is not carried through. Some of the strongest results in the corpus are obtained in split-step simulation environments[18,31] and some on laboratory hardware[19,24,28]. Both are legitimate, and the generalisation risk attached to them is quite different, but earlier tabulations rarely mark which is which.

The consequence is that this literature currently cannot answer its own central question, which is whether one monitoring approach is better than another. That is not a criticism of any individual paper. It is a coordination failure, and coordination failures are fixed by convention rather than by better algorithms. Table 10 sets out the minimum reporting set we propose. Every item in it is information the authors already possess at the time of writing, so the cost of adopting it is close to zero, and the benefit is that the next review of this field will be able to do what this one could not.

**Table 10.** Proposed minimum reporting set for OPM and MFI studies

| # | Item | What to report | What it makes possible |
|---|---|---|---|
| M1 | Evaluation environment | Simulation platform and version, or experimental testbed with component list | Separates simulated from measured results, which is the largest single source of over-optimism |
| M2 | Link specification | Channel count, symbol rate, spacing, span length and count, launch power range, amplifier type | Lets a reader judge whether a claimed accuracy transfers to their own system |
| M3 | Format set | Every modulation format the model was trained and tested on, listed explicitly | Makes classification accuracy interpretable, since difficulty scales with set size |
| M4 | Metric definition | The exact statistic, written as a formula, not only its name | Prevents RMSE, standard error and mean absolute error from being compared as though they were interchangeable |
| M5 | Operating range | The OSNR, dispersion or power range over which the stated accuracy holds | Distinguishes a monitor that works across an operating window from one that works at a point |
| M6 | Data partition | Sizes of the training, validation and test sets, and how they were separated | Exposes leakage between training and test data, the most common failure in learned estimators |
| M7 | Model cost | Parameter count, multiply-accumulate operations per estimate, and inference latency | Turns the deployability question from an assertion into a number, addressing the weakest item in the quality appraisal |
| M8 | Baseline | The analytical or DSP method the learned monitor replaces, evaluated on the same data | Establishes whether learning is necessary, which surprisingly few studies in the corpus demonstrate |

### *3.1.8. Bottlenecks and emerging directions*

The obstacles in this field are consistent across tasks, which suggests they are properties of the problem rather than of particular solutions. Table 11 maps each to its cause, to the evidence for it in the corpus, and to the direction most likely to relieve it.

Three of these deserve emphasis. Labelled data is the binding constraint, not model capacity. Every supervised monitor in the corpus assumes a labelled dataset spanning the operating range, and obtaining one from a production network means deliberately degrading a live service. This is why the automated collection work [20] and the learned channel models [22,25] matter more than their citation counts suggest: a sufficiently faithful differentiable channel model would supply training data without touching the network, and would simultaneously provide the environment that reinforcement learning needs. Generative models are the other candidate route to synthetic training data, and their behaviour and failure modes are documented well enough outside this field [49] for the transfer to be attempted.

Inference cost is a hardware problem, not an algorithmic one. The memory-bandwidth limit on recurrent monitors [18] and the branch-count growth in deep equalisers [19] are both bounded by data movement rather than arithmetic, and no amount of architectural refinement removes them. This is the strongest practical argument for the photonic computing line [10,39,40], which sits at the edge of this corpus but addresses its central constraint directly.

Generalisation is where the field has been least rigorous. A model trained at one symbol rate, on one format set, over one link, and tested on data drawn from the same distribution, tells us little about behaviour after a route change or an equipment upgrade. Transfer learning [30] and physically structured compensators [24] are the two credible responses in the corpus, and they are complementary: the first reduces the cost of adaptation, the second reduces the need for it.

## 4 CONCLUSION

This review shows that AI-assisted optical performance monitoring, modulation format identification, and fibre nonlinearity compensation have evolved predominantly around supervised deep learning, with feed-forward, convolutional, and recurrent models accounting for most primary studies, while reinforcement learning remains an unexplored area in the primary literature. Across the reviewed evidence, input representation emerges as a stronger determinant of practical deployability than network architecture because the computational and memory costs associated with signal transformation can outweigh the cost of the learning model itself. Although individual studies report substantial improvements in monitoring accuracy and transmission performance, their results cannot be reliably ranked or pooled because of differences in metrics, test conditions, modulation-format sets, datasets, and the use of simulation versus experimental evidence. The principal barriers to

**Table 11.** Bottlenecks, causes, evidence in the corpus and emerging directions

| Bottleneck | Root cause | Evidence in the corpus | Emerging direction |
|---|---|---|---|
| Labelled training data | Labels require deliberate degradation of a live link or exhaustive simulation | Slow data collection reported for ROADM switching [34]; dataset collection treated as the research problem in [20] | Differentiable channel models and digital twins [22,25]; automated collection frameworks [20] |
| Training compute | Deep monitors need long training on large datasets | Training reported as resource- and time-intensive [12]; error-free estimation only at high epoch counts [21] | Transfer learning between formats [30]; multi-task sharing of representations [13] |
| Inference latency and memory bandwidth | Data movement, not arithmetic, bounds recurrent and image-based models | Memory-bandwidth limit in LSTM layers [18]; complexity growth with branch count [19]; rendering cost of image representations [11] | Compact representations [14,28]; complexity-aware model selection [32]; optical and photonic computing [10,39,40] |
| Generalisation across conditions | Models are trained and tested on the same distribution | Every primary study evaluates on its own link; no shared benchmark exists | Transfer learning [30]; physically structured compensators that retain model form [24] |
| Comparability of results | No shared metric, dataset or reference link | At least five distinct accuracy metrics across the primary studies (Section 3.1.7) | Adoption of a minimum reporting set (Table 10); community benchmark links |
| Absence of closed-loop control | Monitors estimate but do not act; no safe environment exists for a learning agent | No reinforcement learning study in the primary corpus, despite it being proposed in [1,4–6] | Validated digital twins as a training environment; coupling monitoring output to control-plane decisions [6,44] |

practical deployment are therefore not simply model capacity but the availability of labelled and representative data, memory and inference costs, and generalisation across changing network conditions. To make future progress measurable, the field requires consistent reporting of datasets, experimental conditions, train–test separation, metrics, computational complexity, latency, and generalisation performance through a common minimum reporting framework. More fundamentally, progress toward autonomous optical networks depends on establishing public benchmark datasets, validated differentiable models of deployed fibre links, transparent reporting of inference cost, and closed-loop integration between learned monitoring and network control. Addressing these gaps would shift the field from demonstrating isolated algorithmic improvements toward reproducible, comparable, and operationally meaningful advances in intelligent optical networking.

## 5 DISCLOSURES

The authors declare that there are no funding sources, financial interests, commercial affiliations or other potential conflicts of interest that could have influenced the objectivity of this review or the writing of this paper. No code was developed or used for this study. All records analysed are listed in the references, and the search strings and eligibility criteria needed to reproduce the corpus are given in Section 2.

## References

1) Mata J, et al. Artificial intelligence (AI) methods in optical networks: A comprehensive survey. *Optical Switching and Networking*. 2018;28:43–57. https://doi.org/10.1016/j.osn.2017.12.006.
2) Liu X, et al. AI-based modeling and monitoring techniques for future intelligent elastic optical networks. *Applied Sciences*. 2020;10(1). https://doi.org/10.3390/app10010363.
3) Dong Z, Khan FN, Sui Q, Zhong K, Lu C, Lau APT. Optical performance monitoring: A review of current and future technologies. *Journal of Lightwave Technology*. 2016;34(2):525–543. https://doi.org/10.1109/JLT.2015.2480798.
4) Khan FN, Fan Q, Lu C, Lau APT. An optical communication's perspective on machine learning and its applications. *Journal of Lightwave Technology*. 2019;37(2):493–516. https://doi.org/10.1109/JLT.2019.2897313.
5) Musumeci F, et al. An overview on application of machine learning techniques in optical networks. *IEEE Communications Surveys and Tutorials*. 2019;21(2):1383–1408. https://doi.org/10.1109/COMST.2018.2880039.
6) Rafique D, Velasco L. Machine learning for network automation: Overview, architecture, and applications [Invited Tutorial]. *Journal of Optical Communications and Networking*. 2018;10(10):D126–D143. https://doi.org/10.1364/JOCN.10.00D126.
7) Wang D, Zhang M. Artificial intelligence in optical communications: From machine learning to deep learning. *Frontiers in Communications and Networks*. 2021;2. https://doi.org/10.3389/frcmn.2021.656786.
8) Khan FN, Fan Q, Lu C, Lau APT. Machine learning-assisted optical performance monitoring in fiber-optic networks. *IEEE Photonics Society Summer Topicals Meeting Series (SUM)*. 2018;p. 55–56. https://doi.org/10.1109/PHOSST.2018.8456700.

9) Wei J, Yi L, Giacoumidis E, Cheng Q, Lau APT. Special issue on 'optics for AI and AI for optics'. *Applied Sciences*. 2020;10(9). https://doi.org/10.3390/app10093262.
10) Gigan S, Krzakala F, Daudet L, Carron I. Artificial intelligence: From electronics to optics. *Photoniques*. 2020;(104):49–52. https://doi.org/10.1051/photon/202010449.
11) Wang D, et al. Intelligent constellation diagram analyzer using convolutional neural network-based deep learning. *Optics Express*. 2017;25(15):17150. https://doi.org/10.1364/OE.25.017150.
12) Khan FN, et al. Joint OSNR monitoring and modulation format identification in digital coherent receivers using deep neural networks. *Optics Express*. 2017;25(15):17767. https://doi.org/10.1364/OE.25.017767.
13) Fan X, et al. Joint optical performance monitoring and modulation format/bit-rate identification by CNN-based multi-task learning. *IEEE Photonics Journal*. 2018;10(5). https://doi.org/10.1109/JPHOT.2018.2869972.
14) Yi A, et al. Modulation format identification and OSNR monitoring using density distributions in Stokes axes for digital coherent receivers. *Optics Express*. 2019;27(4):4471. https://doi.org/10.1364/OE.27.004471.
15) Amari A, Lin X, Dobre OA, Venkatesan R, Alvarado A. A machine learning-based detection technique for optical fiber nonlinearity mitigation. *IEEE Photonics Technology Letters*. 2019;31(8):627–630. https://doi.org/10.1109/LPT.2019.2902973.
16) Karanov B, et al. End-to-end deep learning of optical fiber communications. *Journal of Lightwave Technology*. 2018;36(20):4843–4855. https://doi.org/10.1109/JLT.2018.2865109.
17) Saif WS, Esmail MA, Ragheb AM, Alshawi TA, Alshebeili SA. Machine learning techniques for optical performance monitoring and modulation format identification: A survey. *IEEE Communications Surveys and Tutorials*. 2020;22(4):2839–2882. https://doi.org/10.1109/COMST.2020.3018494.
18) Wang Z, Yang A, Guo P, He P. OSNR and nonlinear noise power estimation for optical fiber communication systems using LSTM based deep learning technique. *Optics Express*. 2018;26(16):21346. https://doi.org/10.1364/OE.26.021346.
19) Lu X, et al. Memory-controlled deep LSTM neural network post-equalizer used in high-speed PAM VLC system. *Optics Express*. 2019;27(5):7822. https://doi.org/10.1364/OE.27.007822.
20) Shao C, Giacoumidis E, Billah SM, et al. Machine Learning in Short-Reach Optical Systems: A Comprehensive Survey. *Photonics*. 2024;11(7):613. https://doi.org/10.3390/photonics11070613.
21) Wang D, et al. Modulation format recognition and OSNR estimation using CNN-based deep learning. *IEEE Photonics Technology Letters*. 2017;29(19):1667–1670. https://doi.org/10.1109/LPT.2017.2742553.
22) Ma Z, Jia P, Han D, Zhang M, Ghassemlooy Z, Wang L. Deep-learning-based channel estimation for multi-wavelength visible light communication system. *West Asian Symposium on Optical and Millimeter-Wave Wireless Communications (WASOWC)*. 2022;p. 1–4. https://doi.org/10.1109/WASOWC54657.2022.9798444.
23) Hanawa M, Goto Y. Exploration of sub-optimal constellation maps for coherent fiber-optic communication links by auto-encoder type machine learning. *International Electrical Engineering Congress (iEECON)*. 2022. https://doi.org/10.1109/iEECON53204.2022.9741572.
24) Oliari V, et al. Revisiting efficient multi-step nonlinearity compensation with machine learning: An experimental demonstration. *Journal of Lightwave Technology*. 2020;38(12):3114–3124. https://doi.org/10.1109/JLT.2020.2994220.
25) Wang D, et al. Data-driven optical fiber channel modeling: A deep learning approach. *Journal of Lightwave Technology*. 2020;38(17):4730–4743. https://doi.org/10.1109/JLT.2020.2993271.
26) Tschannen M, Bachem O, Lucic M. Recent advances in autoencoder-based representation learning. *arXiv preprint*. 2018. https://arxiv.org/abs/1812.05069.
27) Doan VS, Huynh-The T, Hua CH, Pham QV, Kim DS. Learning constellation map with deep CNN for accurate modulation recognition. *arXiv preprint*. 2020. https://arxiv.org/abs/2009.02026.
28) Locatelli F, Christodoulopoulos K, Moreolo MS, Fabrega JM, Spadaro S. Machine learning-based in-band OSNR estimation from optical spectra. *IEEE Photonics Technology Letters*. 2019;31(24):1929–1932. https://doi.org/10.1109/LPT.2019.2950058.
29) Wang D, et al. Machine learning-based multifunctional optical spectrum analysis technique. *IEEE Access*. 2019;7:19726–19737. https://doi.org/10.1109/ACCESS.2019.2895409.
30) Zhang J, et al. Joint modulation format identification and OSNR monitoring using cascaded neural network with transfer learning. *IEEE Photonics Journal*. 2021;13(1). https://doi.org/10.1109/JPHOT.2021.3056471.
31) Zhao Y, et al. Low-complexity fiber nonlinearity impairments compensation enabled by simple recurrent neural network with time memory. *IEEE Access*. 2020;8:160995–161004. https://doi.org/10.1109/ACCESS.2020.3021146.
32) Deligiannidis S, Mesaritakis C, Bogris A. Performance and complexity evaluation of recurrent neural network models for fibre nonlinear equalization in digital coherent systems. *European Conference on Optical Communications (ECOC)*. 2020. https://doi.org/10.1109/ECOC48923.2020.9333288.
33) Shahkarami A, Yousefi MI, Jaouen Y. Efficient deep learning of nonlinear fiber-optic communications using a convolutional recurrent neural network. *20th IEEE International Conference on Machine Learning and Applications (ICMLA)*. 2021;p. 668–673. https://doi.org/10.1109/ICMLA52953.2021.00112.
34) Mo W, Gutterman CL, Li Y, Zhu S, Zussman G, Kilper DC. Deep-neural-network-based wavelength selection and switching in ROADM systems. *Journal of Optical Communications and Networking*. 2018;10(10). https://doi.org/10.1364/JOCN.10.0000D1.
35) Lun H, et al. Soft failure identification for long-haul optical communication systems based on one-dimensional convolutional neural network. *Journal of Lightwave Technology*. 2020;38(11):2992–2999. https://doi.org/10.1109/JLT.2020.2989153.
36) Karanov B, Chagnon M, Aref V, Lavery D, Bayvel P, Schmalen L. Optical fiber communication systems based on end-to-end deep learning (Invited Paper). *IEEE Photonics Conference (IPC)*. 2020. https://doi.org/10.1109/IPC47351.2020.9252544.
37) Zhang Q, Wang Z, Duan S, Jiang N, Cao B, Wu Y. An improved end-to-end optical transmission system based on deep learning. *19th International Conference on Optical Communications and Networks (ICOCN)*. 2021. https://doi.org/10.1109/ICOCN53177.2021.9563723.
38) Neskorniuk V, et al. End-to-end deep learning of long-haul coherent optical fiber communications via regular perturbation model. *European Conference on Optical Communication (ECOC)*. 2021. https://doi.org/10.1109/ECOC52684.2021.9605928.
39) Sadeghzadeh H, Koohi S, Paranj AF. Free-space optical neural network based on optical nonlinearity and pooling operations. *IEEE Access*. 2021;9:146533–146549. https://doi.org/10.1109/ACCESS.2021.3123230.
40) Piccinotti D, MacDonald KF, Gregory SA, Youngs I, Zheludev NI. Artificial intelligence for photonics and photonic materials. *Reports on Progress in Physics*. 2021;84(1). https://doi.org/10.1088/1361-6633/abb4c7.
41) Song Y, Wang D, Ye H, Qin J, Zhang M. Wavelength controllable forward prediction and inverse design of nanophotonic devices using deep learning. *European Conference on Optical Communications (ECOC)*. 2020. https://doi.org/10.1109/ECOC48923.2020.9333397.

42) Amirabadi MA, Kahaei MH, Nezamalhosseini SA. A deep learning based detector for FSO system considering imperfect CSI scenario. *3rd West Asian Symposium on Optical and Millimeter-Wave Wireless Communications (WASOWC)*. 2020. https://doi.org/10.1109/WASOWC49739.2020.9410006.
43) Zhang G, Wang Y, Ding H, Gao W, Gao K, Ma G. Deep learning based optical network layer recovery mechanism for critical services of power communication network. *19th International Conference on Optical Communications and Networks (ICOCN)*. 2021. https://doi.org/10.1109/ICOCN53177.2021.9563732.
44) Furdek M, Natalino C. Machine learning for optical network security management. *Optical Fiber Communication Conference (OFC)*. 2020. Paper M4E.4. https://doi.org/10.1364/OFC.2020.M4E.4.
45) Morais RM, Pedro J. Evaluating machine learning models for QoT estimation. *International Conference on Transparent Optical Networks (ICTON)*. 2018;2018-July. https://doi.org/10.1109/ICTON.2018.8473941.
46) Rafique D. Machine learning based optimal modulation format prediction for physical layer network planning. *International Conference on Transparent Optical Networks (ICTON)*. 2018;2018-July. https://doi.org/10.1109/ICTON.2018.8473593.
47) Devi PP, Vincent, Simatupang JW. Machine learning-based modulation format identification and optical performance monitoring techniques implementation. *Advanced Communication Technologies and Signal Processing (ACTS)*. 2021. https://doi.org/10.1109/ACTS53447.2021.9708283.
48) Jha DK, Mishra JK. Performance evaluation of DWDM optical transmission system using deep learning technique. *IEEE Photonics Conference (IPC)*. 2021. https://doi.org/10.1109/IPC48725.2021.9592951.
49) Gui J, Sun Z, Wen Y, Tao D, Ye J. A review on generative adversarial networks: Algorithms, theory, and applications. *arXiv preprint*. 2020. https://arxiv.org/abs/2001.06937.