

Transfer Learning Approaches for Cross-Platform Intelligent Fault Prediction

Sundararajan G^{1*}, Ramanathan S² and Ananthakrishnan R³

¹Department of Computational Intelligence, Tamiraparani Center for Cognitive Computing, Indonesia Email: ganesh.sundararajan@tamiraparani.edu

ORCID: <https://orcid.org/7542-9994-7843-6318>

²Department of Robotics and Automation, Noyyal Center for Future Computing Architectures, South Korea Email: suresh.ramanathan@noyyal.edu

ORCID: <https://orcid.org/5571-0744-5388-1613>

³Department of Robotics and Automation, Bharathapuzha Center for Neural Computation, Sri Lanka Email: rajesh.ananthakrishnan@bharathapuzha.edu

ORCID: <https://orcid.org/7906-0244-9607-9939>

ABSTRACT

Deep learning has pushed intelligent fault diagnosis to near-perfect accuracy on the laboratory test rigs where it is usually trained and evaluated, yet that accuracy degrades sharply the moment a model is asked to work on a different machine, a different operating condition, or a different platform in the same industrial fleet -- the well-known problem of domain shift. Transfer learning and domain adaptation address this gap directly by reusing knowledge learned on a data-rich source domain to build working diagnostic models on target machines where labeled fault data is scarce or absent. This article synthesizes the literature on cross-platform transfer learning for intelligent fault prediction, tracing the field from its foundations in transfer learning theory and domain-adaptation generalization bounds, through deep convolutional and autoencoder architectures purpose-built for vibration-signal diagnostics, to subdomain-level and adversarial alignment methods that correct the class-collapse failure of naive global distribution matching, and on to harder, more industrially realistic settings: few-shot and meta-learning-based cross-machine transfer, open-set recognition of previously unseen fault modes, and federated transfer learning across organizational boundaries. We give particular weight to the industrial IoT and edge-computing literature that determines whether any of this can actually run on the constrained hardware co-located with real equipment. Read together, the evidence supports transfer learning as a mature and repeatedly validated technique for closing much of the cross-domain accuracy gap in fault diagnosis, while leaving multi-source fleet-scale transfer, open-set safety, and edge-deployable accuracy as an open, still largely benchmark-rig-bound research agenda.

Keywords: transfer learning; fault diagnosis; domain adaptation; industrial Internet of Things; cross-platform prediction

Citation: Sundararajan G, Ramanathan S and Ananthakrishnan R [2023]. Transfer Learning Approaches for Cross-Platform Intelligent Fault Prediction. DOI: <https://doi.org/10.5281/zenodo.22843870>

Copyright: © 2023 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 9 September 2022 Accepted: 14 December 2022 Published: 14 March 2023

Article Class: Research Article

1. Introduction

Industrial equipment fails in ways that are expensive to predict and even more expensive to get wrong. A bearing that seizes on an offshore wind turbine, a gearbox that strips its teeth in a rolling mill, or a motor that overheats on a factory floor each carries a cost that is measured not just in the part itself but in unplanned downtime, safety exposure, and cascading damage to adjacent machinery. Intelligent fault diagnosis -- using sensor data (vibration, current, acoustic emission, temperature) together with machine learning to detect, classify, and eventually predict incipient faults -- has therefore been one of the most consistently funded applications of pattern recognition in industrial engineering for more than two decades. Deep learning has, over the past ten years, pushed the state of the art in this area substantially: convolutional and recurrent architectures trained on raw or lightly processed vibration signals now routinely outperform classical approaches built on hand-crafted features and shallow classifiers, and the numbers reported in benchmark studies are, by the standards of industrial diagnostics, remarkably good -- diagnostic accuracies in the high nineties on standard bearing datasets are common [1,2].

The trouble is that those numbers describe a narrower situation than the phrase "intelligent fault diagnosis" suggests. A deep network trained to recognize inner-race, outer-race, and ball-fault vibration signatures on one motor, at one rotational speed, under one load, on one test rig, is a model of that motor under those conditions -- not a general-purpose fault detector. Move the same sensor to a different machine of a different make, change the operating speed, swap in a different bearing geometry, or simply move from a laboratory test rig to a real production line, and the underlying vibration statistics shift enough that the trained classifier's accuracy can fall sharply, sometimes to little better than chance on the fault classes it was never shown in the new regime. This is the well-known problem of *domain shift*: the assumption that training and deployment data are drawn from the same distribution, an assumption nearly every supervised learning method depends on implicitly, is routinely violated the moment an intelligent fault diagnosis system leaves the laboratory and is asked to work across the heterogeneous population of machines, sensors, and operating regimes that make up a real industrial fleet or a *platform* of physically distinct but functionally related

equipment.

Transfer learning is the branch of machine learning built specifically to address this gap: rather than assuming source and target data share a distribution, transfer learning asks how knowledge learned on a data-rich *source domain* can be reused, adapted, or partially retained when only limited or unlabeled data is available in a related *target domain* [3]. Applied to fault diagnosis, the promise is direct and practically important: train a diagnostic model where labeled fault data is abundant -- a well-instrumented test rig, a fleet of machines with a long failure history, a public benchmark dataset -- and transfer that model, with modest or no additional labeled data, to a new machine, a new operating condition, or a new sensor placement where labeled fault examples are scarce or entirely absent, which is the normal condition for a newly commissioned asset or a machine that has, fortunately, not yet failed enough times to build a labeled fault dataset of its own. This article surveys the literature that has grown up around exactly that problem: cross-domain and cross-platform transfer learning for intelligent fault prediction, spanning domain adaptation theory, deep transfer architectures purpose-built for vibration-based diagnostics, subdomain and adversarial adaptation methods, multi-source and few-shot variants, and the industrial Internet of Things (IIoT) and edge-computing constraints that determine whether any of this actually runs where the sensors are.

Although this synthesis concentrates on discrete fault classification, the same domain-shift problem, and much of the same transfer-learning machinery, applies equally to the closely related task of remaining-useful-life (RUL) prediction, where Zhang, Wang, Li, Cui, Liu, Yang, and Hu showed that a deep recurrent network's RUL estimates, learned on one operating condition, transfer usefully to a different condition once fine-tuned with only limited target data, rather than needing an entirely new model trained from scratch for every new regime [4]. We use "cross-platform" deliberately broadly, in the way the literature itself does: it covers transfer across different physical machines of the same type (motor A to motor B), transfer across operating conditions on the same machine (one load or speed regime to another), transfer across different but related machine types (bearings to gearboxes, or one turbine model to another), and

transfer across sensor modalities or placements. What unifies these settings, and what justifies treating them under a single research umbrella, is that in every case the underlying physics of fault initiation and propagation is similar enough that a model trained on one instance genuinely does carry useful information about the others -- the task is not to learn something new from scratch, but to identify and reuse the invariant structure while adapting to what differs. This article is organized as follows. We first lay out the foundations of transfer learning and domain adaptation as a formal problem. We then trace how the fault diagnosis community adapted general-purpose transfer learning machinery -- feature-alignment methods, adversarial domain adaptation, subdomain adaptation -- to the specific structure of vibration and condition-monitoring signals. We follow with a section on the harder, more realistic settings that recent work has pushed into: multi-source and cross-machine transfer, wind-turbine and gearbox case studies, few-shot and open-set diagnosis. We then turn to the deployment side of the problem -- how any of this is supposed to run on real industrial IoT hardware and communicate across a fleet -- before closing with a candid look at what remains unresolved.

It is worth being explicit about the nature of the evidence base we are synthesizing. The great majority of the results reported in this literature come from public or semi-public benchmark datasets -- test rigs at Case Western Reserve University, Paderborn University, Jiangnan University, and comparable university laboratories -- rather than from long-running deployments on production equipment, and cross-domain transfer is usually evaluated between artificially constructed "domains" (different load levels on the same rig, for instance) rather than between genuinely different machines in the field. That is a real limitation, and one the field is only gradually correcting as industrial partners publish field data. What the literature supports with real confidence is that the transfer learning techniques developed here reliably improve target-domain diagnostic performance relative to a model trained only on source data and tested cold on the target, and that the gap they close is often large -- tens of percentage points of accuracy in the harder cross-condition settings. Whether those gains hold up at the scale and heterogeneity of a real industrial platform with hundreds of machines is a question this synthesis

returns to explicitly in its discussion of open challenges.

2. Foundations of Transfer Learning and Domain Adaptation

Pan and Yang's widely cited 2010 survey gave the field its now-standard vocabulary: a *domain* is defined by a feature space and a marginal probability distribution over that space, and a *task* is defined by a label space and an objective (predictive) function; transfer learning is then the general problem of improving a target task's learning using knowledge from a different but related source domain and/or task [3]. Within that framework, several distinctions matter a great deal for fault diagnosis specifically. *Homogeneous* transfer assumes source and target share the same feature space (the same sensor type and sampling configuration, for instance) but differ in marginal or conditional distribution -- this is the dominant case in vibration-based fault diagnosis, where the same accelerometer signal is collected under different conditions. *Heterogeneous* transfer, by contrast, allows the feature spaces themselves to differ, which becomes relevant when transferring between, say, vibration and acoustic emission sensors, or between machines instrumented with genuinely different sensor suites -- a harder and less thoroughly solved problem in this literature. *Inductive* transfer assumes some labeled target data is available; *transductive* transfer (the setting usually called unsupervised domain adaptation) assumes none, only unlabeled target data, which is the setting closest to the real deployment problem of a newly installed machine with no fault history of its own.

The formal justification for why domain adaptation should be expected to work at all, rather than being an ad hoc engineering trick, comes from learning theory. Ben-David and colleagues established a generalization bound for domain adaptation showing that a classifier's expected error on a target domain is bounded by three terms: its error on the source domain, a measure of the divergence between the source and target marginal distributions (the H-divergence, computable from how well a classifier can distinguish source from target examples), and a term capturing how well the ideal joint hypothesis performs on both domains combined [5]. This bound is the theoretical backbone of essentially every domain-adaptation method used in fault diagnosis: it says explicitly that reducing the measurable distributional gap between source and target feature

representations -- rather than simply hoping a source-trained model generalizes -- is what should improve target performance, and it is precisely this insight that motivated the shift, described in the next section, from statistical feature-alignment techniques to explicitly adversarial ones that train a representation to make source and target indistinguishable.

Two families of technique operationalize that insight. The first, statistical distribution matching, minimizes an explicit distance between source and target feature distributions during training. Sun, Feng, and Saenko's CORAL (correlation alignment) approach is a particularly simple and influential instance: rather than any deep architectural change, CORAL re-aligns the second-order statistics (the feature covariance) of the source and target features by a linear transformation, and the authors showed this "frustratingly easy" recoloring step recovers much of the benefit of far more complex domain-adaptation pipelines on standard visual-domain-adaptation benchmarks [6]. Long, Cao, Cao, Wang, and Jordan's Deep Adaptation Network (DAN) line of work takes the same statistical-matching idea into deep architectures directly, embedding a multi-kernel maximum mean discrepancy (MK-MMD) penalty into intermediate layers of a deep network so that the learned representation itself is trained to minimize the distributional gap between domains, rather than aligning hand-designed features after the fact [7]. The second family, adversarial domain adaptation, borrows the generative-adversarial-network idea directly: a feature extractor is trained jointly with a domain discriminator whose job is to tell source features from target features, and the feature extractor is trained, via a gradient-reversal or minimax objective, to fool that discriminator -- the practical effect being that the learned representation becomes progressively less informative about which domain an example came from, while (with an appropriate label-prediction loss retained on the source) remaining informative about the fault class. This general architecture, and the substantial family of adversarial and statistical-matching variants it spawned, is the shared technical toolkit that the fault-diagnosis-specific literature discussed below draws on almost universally.

3. Deep Transfer Learning Architectures for Fault Diagnosis

The fault diagnosis community began adapting

general-purpose transfer learning machinery to vibration-signal diagnostics in earnest around 2017-2019, and a handful of architectural patterns established in that period remain the backbone of the field. Lu, Liang, Cheng, Meng, Yang, and Zhang's early deep-model domain adaptation work is a useful starting point: they trained a deep neural feature extractor with an MMD-style domain-discrepancy penalty added to the standard classification loss, directly following the DAN-style statistical-matching approach, and evaluated the resulting model on cross-load transfer between different operating conditions of the same rotating machinery, reporting substantially higher target-domain diagnostic accuracy than a source-only baseline transferred without adaptation [8]. Wen, Gao, and Li pushed the same idea through a sparse autoencoder architecture rather than a purely discriminative one, using the autoencoder's reconstruction objective together with an MMD-based domain-adaptation term to learn a fault-discriminative representation that additionally aligns source and target feature distributions, and reported strong transfer accuracy across several fault-diagnosis transfer tasks constructed from standard bearing and gearbox datasets [9].

Guo, Lei, Xing, Yan, and Li's Deep Convolutional Transfer Learning Network (DCTLN) is one of the more influential architectures in this specific lineage, explicitly designed around the two-stage structure that recurs throughout the field: a condition-recognition stage that trains a convolutional feature extractor to distinguish fault types on labeled source data, followed by a domain-adaptation stage in which the same feature extractor is further tuned, using a domain-discrepancy loss computed between source and unlabeled target features, so that the same convolutional filters that discriminate fault classes on the source domain also produce domain-invariant representations on the target [10]. The paper's framing of the problem -- one labeled, data-rich source condition and one essentially unlabeled target condition, with the goal of building a working diagnostic classifier for the target using no target labels at all -- captures precisely the deployment scenario that makes transfer learning practically valuable for fault diagnosis: a newly commissioned machine, or a machine operating under a new load regime, very often has plenty of *unlabeled* vibration data (the sensor is already recording) but essentially no labeled fault data, because faults are rare events and a

machine that has not yet failed cannot supply examples of its own failure modes.

Zhao, Zhang, Zhan, and Pang extended the convolutional transfer-network idea to a multi-scale architecture, arguing that fault-discriminative information in vibration signals is distributed across multiple time-frequency scales simultaneously (a bearing defect produces both broadband impulsive energy and narrowband resonance effects at specific frequencies related to defect geometry and rotational speed) and that a single-scale convolutional feature extractor, of the kind used in earlier transfer architectures, under-exploits this structure; their Deep Multi-Scale Convolutional Transfer Learning Network combines parallel convolutional branches at different receptive-field scales with a domain-adaptation loss applied across the fused multi-scale representation, and reported improved cross-domain diagnostic accuracy relative to single-scale transfer baselines across several transfer tasks built on standard bearing datasets [11]. Taken together, this line of work established a design pattern that later papers, including several discussed in the next section, largely inherit: a convolutional or autoencoder-based feature extractor tuned for vibration-signal structure specifically, combined with an explicit domain-discrepancy penalty (MMD-style or adversarial) added to an otherwise standard fault-classification training objective, evaluated on cross-load or cross-machine transfer tasks constructed from a shared public benchmark dataset.

4. Subdomain Adaptation, Adversarial Methods, and Open-Set Diagnosis

A limitation of the global-distribution-matching approaches described above is precisely that they are global: minimizing the discrepancy between the overall source and target feature distributions does not guarantee that individual fault classes align correctly with each other across domains, and it is entirely possible for a global-alignment method to produce a representation where the source and target marginal distributions look similar in aggregate while individual class-conditional distributions -- inner-race faults in the source mapping near outer-race faults in the target, for instance -- are badly misaligned. Zhu and colleagues' Deep Subdomain Adaptation Network (DSAN) addressed this directly by replacing the global MMD penalty with a local, class-wise MMD computed

per-class using pseudo-labels for the (unlabeled) target domain, so that alignment is enforced between corresponding fault subclasses rather than between the domains as undifferentiated wholes; the original paper demonstrated this subdomain-level alignment improves classification accuracy over global-alignment methods on standard domain-adaptation benchmarks, and the subdomain-adaptation principle it established has since been adopted extensively within fault diagnosis specifically [12]. Wang, Zhang, Xu, Wang, and Yang's cross-domain fault diagnosis method built directly on this subdomain-adaptation logic, combining it with a few-shot learning setup in which only a handful of labeled target examples are assumed available per fault class, and reported that subdomain-level alignment produced meaningfully better few-shot cross-domain diagnostic performance than global-alignment baselines on rotating-machinery transfer tasks [13].




Figure 1: Average diagnostic F1-score of Lei et al.'s open-set domain-adaptation method, evaluated separately on three independent rolling-bearing benchmark datasets (CWRU, SEU, PU), showing that the joint known-fault-recognition-plus-unknown-fault-detection objective holds up consistently across datasets rather than being an artifact of a single benchmark [18].

Adversarial domain adaptation has become at least as prominent as statistical-matching approaches within fault diagnosis specifically, generally by combining a domain-adversarial training objective with an architecture adapted to vibration-signal or graph-structured sensor data. Li, Jin, Xu, and Yang's domain-adversarial multi-graph convolutional network is a representative recent example: rather than treating vibration samples as independent feature vectors, they construct a graph representation connecting sensor channels or feature nodes and apply graph convolution to capture structural relationships among them, combined with a domain-adversarial loss (a gradient-reversal domain discriminator in the now-standard style) so that the learned graph representation becomes invariant to the source-versus-target distinction while remaining

discriminative for fault type; the authors reported strong diagnostic accuracy on unsupervised cross-condition transfer tasks relative to non-adversarial and non-graph baselines [14]. Javed, Ding, Yan, Ma, and Kumar's more recent explainable domain-adaptive CNN-Transformer combines convolutional local feature extraction with a Transformer's self-attention mechanism for longer-range temporal dependence in vibration signals, applies domain adaptation in both source-to-target and target-to-source directions (bidirectional adaptation, intended to make the learned alignment more symmetric and robust than single-direction methods), and adds interpretability analysis -- attention visualization tied to physically meaningful fault-frequency components -- that most earlier transfer architectures did not attempt, addressing a practical objection that a "black box" domain-adapted classifier is a harder sell to maintenance engineers who need to trust and act on its output [15]. Wang, Lv, Chen, Zhang, and Li's maximum-margin local domain adaptation method pursues a related goal of finer-grained, class-aware alignment by explicitly maximizing the margin between different fault classes' local neighborhoods across domains rather than only minimizing a global or per-class discrepancy term, reporting improved diagnostic accuracy under multiple operating conditions relative to methods that align distributions without an explicit margin objective [16]. Jiao, Zhou, Liu, Ma, and Xu's lightweight domain-adaptation architecture for defected-bearing diagnosis takes the opposite emphasis, prioritizing a small parameter count and low inference cost alongside domain adaptation specifically because, as we discuss further below, most of the architectures in this section were designed and evaluated without much attention to whether they could run on constrained industrial hardware at all [17].

A separate but related strand addresses *open-set* diagnosis, a setting the standard closed-set domain-adaptation formulation does not handle: a target machine may develop a fault mode that never appeared in the labeled source data at all, and a closed-set classifier, adversarially trained or otherwise, will confidently -- and wrongly -- assign such a novel fault to one of its known classes rather than flagging it as something new. Lei, Pan, Hu, He, and Li's open-set domain-adaptation method for rolling-bearing diagnosis combines adversarial feature alignment for the known, shared fault classes with an explicit mechanism

for recognizing and separating out target samples that do not belong to any known source class, reporting that their method both preserves strong closed-set transfer accuracy and substantially improves the ability to correctly flag genuinely novel target-domain fault patterns relative to closed-set adversarial baselines applied naively to the open-set problem [18]. This open-set capability matters a great deal for real deployment: a platform-wide fault-prediction system that has been trained, even very well, on the fault modes documented in a source fleet's history will eventually encounter a failure mode nobody has seen before, and a system that cannot represent "I don't recognize this" is a liability rather than a safeguard in exactly the moment it matters most.

5. Cross-Machine, Multi-Source, and Few-Shot Transfer

Most of the architectures discussed so far are evaluated on cross-condition transfer within a single test rig -- different loads or speeds on the same physical machine -- which is a genuinely useful and well-controlled experimental setting but is not the same problem as transferring across a heterogeneous fleet of *different* machines, which is closer to what "cross-platform" means in a real industrial context. A smaller but growing body of work targets that harder setting directly. Zhang, Han, Tian, and Shi's domain-adaptation meta-learning network combines a meta-learning outer loop -- learning an initialization or update rule that generalizes across a distribution of related few-shot diagnosis tasks, in the spirit of model-agnostic meta-learning -- with a domain-adaptation inner mechanism and an explicit "discard-supplement" module designed to identify and down-weight source features that do not transfer usefully to a new target machine while supplementing with target-specific structure learned from the available few-shot target examples; the authors reported that this combination outperformed both plain domain adaptation and plain meta-learning baselines on few-shot cross-domain rotating-machinery diagnosis tasks constructed to simulate transfer across genuinely different machine configurations [19]. This meta-learning-plus-domain-adaptation combination is a natural response to a real tension in cross-platform transfer: domain adaptation methods generally assume the target domain, whatever its differences from the source, is fixed and knowable at training time, whereas a platform with many heterogeneous machines effectively

presents a *distribution* of target domains, which is closer to the problem meta-learning is designed to solve.

Table 1: Diagnostic accuracy under mixed operating conditions for a hydrodynamic-transmission bearing fault diagnosis task, comparing a non-transfer baseline, single-technique baselines, and the combined few-shot-learning-plus-transfer-learning (FSL+TL) approach, with and without an added attention mechanism (AM), as reported by Sun et al. [20].

Method	Description	Accuracy (mixed conditions)
SVM	Classical shallow classifier, no transfer	16.19%
WDCNN	Deep CNN trained from scratch, no transfer	69.58%
FSL	Few-shot learning alone, no transfer component	79.71%
WDCNN + TL	Deep CNN with transfer learning, no few-shot component	82.15%
FSL + TL (proposed)	Combined few-shot and transfer learning	85.30%
FSL + TL + AM (proposed)	Combined few-shot and transfer learning with attention mechanism	88.75%

Sun, Yang, and Yang's bearing fault diagnosis method for hydrodynamic transmissions, published in 2025, is a useful recent case study precisely because it targets a specific, named class of industrial equipment (hydrodynamic transmission systems, used in heavy machinery drivetrains) rather than a generic benchmark rig, integrating few-shot learning with transfer learning so that the diagnostic model can be adapted to a new transmission unit using only a small number of labeled fault examples from that unit, combined with knowledge transferred from a richer source dataset [20]. Guo, Wu, Zhang, Long, Chen, Cabrera, and Li's generative transfer learning approach for wind-turbine gearbox diagnosis addresses a version of the cross-platform problem that is genuinely industrially important: wind farms consist of many nominally identical turbines whose gearboxes nonetheless age, and fail, somewhat differently depending on site-specific load history, and a diagnostic model trained on one turbine's fault history (or on bench-test data) needs to transfer to sibling turbines with only limited labeled fault data of their own; their approach uses a generative model to synthesize additional target-like training examples alongside a transfer-learning classifier, and the authors reported improved diagnostic accuracy on gearbox fault classification relative to non-generative transfer baselines, evaluated on wind-turbine gearbox vibration data [21]. An, Ai, and Liu's domain-alignment-and-discriminative-feature-learning

model, evaluated on bearing fault diagnosis, pursues a related but architecturally distinct combination -- explicit domain alignment paired with a discriminative loss term designed specifically to keep fault classes well-separated even as the domain-alignment objective pulls source and target representations together, addressing the same class-collapse risk that motivated the subdomain-adaptation work discussed above from a different angle -- and reported strong cross-condition diagnostic performance on standard bearing transfer tasks [22].

Multi-source transfer -- combining information from several distinct labeled source domains rather than a single one, which is closer still to the situation of a real industrial platform with a diverse installed base -- remains comparatively less developed within fault diagnosis specifically than the single-source-to-single-target setting that dominates the literature surveyed above, though the broader transfer learning and domain adaptation literature, including Long et al.'s deep adaptation network work discussed earlier, provides much of the statistical-matching machinery that multi-source extensions build on directly by aligning each source separately, or jointly, with the shared target [7]. This remains, in our reading, one of the clearer gaps between what has been most thoroughly studied (single source, single target, same equipment family) and what a genuinely cross-platform deployment would actually need (many heterogeneous sources, informative weighting of which source machines are most relevant to a given target, and graceful degradation when no single source is a close match) -- a point we return to in the discussion of open challenges below.

6. Industrial IoT Deployment and Edge Constraints

None of the transfer and domain-adaptation architectures discussed above matter for cross-platform fault prediction in practice unless they can actually be deployed on the sensing and computing infrastructure that real industrial platforms use, and that infrastructure looks considerably more constrained, and more distributed, than the workstation-class GPUs most of this literature is trained and evaluated on. The industrial IoT (IIoT) literature on condition monitoring and fault diagnosis has developed largely in parallel with the deep-transfer-learning literature surveyed above, and

the two need to be read together to assess whether cross-platform fault prediction, as described in the papers above, is realistically deployable. Xu, Zhou, and Fang's work on intelligent roller-bearing fault diagnosis situated explicitly within an Industrial Internet of Things architecture is a useful anchor here, framing the diagnostic model not as a standalone classifier but as one component of a distributed sensing-and-inference pipeline with the communication and reliability constraints that implies for a fleet of networked sensors reporting condition data across an industrial IoT network [23].

Two design responses to these constraints recur across the IIoT fault-diagnosis literature. The first pushes inference, and sometimes training, out to the network edge rather than relying on a central cloud service to process every sensor's data stream, both to reduce the latency between a fault signature appearing and an alarm being raised and to reduce the bandwidth cost of streaming raw high-frequency vibration data continuously off every machine in a fleet. Liu, Zhang, Wu, and Fu's edge-PLC-based fault-diagnosis architecture is a direct instance of this pattern, embedding a random-forest diagnostic model directly on programmable-logic-controller-class edge hardware co-located with the monitored equipment, reporting that this edge placement substantially reduced diagnostic latency relative to a cloud-centralized alternative while maintaining competitive diagnostic accuracy on their industrial IoT test setup [24]. The second response compresses the model itself so that the more capable deep architectures described in earlier sections -- which were generally designed and benchmarked without much regard for parameter count or inference cost -- become small enough to run on that edge hardware in the first place. Wang, Yu, Wu, Wang, Zhou, and Hu's adaptive knowledge-distillation framework for lightweight intelligent fault diagnosis in IoT edge computing follows the now-standard teacher-student distillation pattern, training a larger, more accurate "teacher" diagnostic model centrally and then distilling its behavior into a much smaller "student" network sized for edge deployment, with an adaptive component that adjusts the distillation process to the target edge device's resource profile; the authors reported that the resulting compact student models retained diagnostic accuracy close to their larger teacher counterparts while meeting substantially tighter

memory and compute budgets, evaluated on IoT edge hardware [25]. Jiao et al.'s lightweight domain-adaptation architecture, discussed in the previous section, is squarely in this same tradition, treating a small parameter footprint as a first-class design constraint alongside domain-adaptation accuracy rather than as an afterthought applied to a model designed without it [17].

Privacy and data-sharing constraints add a further wrinkle that is specific to genuinely cross-*platform* deployment, as distinct from cross-condition transfer on a single organization's own equipment: transferring a diagnostic model's knowledge across machines owned or operated by different organizations -- different plants within a conglomerate with separate IT boundaries, or genuinely different companies operating similar equipment who might benefit from pooling diagnostic knowledge without pooling raw sensor data -- runs into exactly the data-governance concerns that motivated federated learning in other domains. Xue, Zhao, Zhang, Ma, Bu, and Peng's federated transfer learning framework with consensus knowledge distillation addresses this directly for fault diagnosis, combining a federated training loop, in which no raw sensor data leaves any participating site, with a knowledge-distillation-based consensus mechanism for reconciling what different sites' locally trained models have learned, and transfer-learning components to handle the fact that different sites' equipment and operating conditions differ; the authors reported that this combination achieved diagnostic accuracy competitive with a centralized-data baseline while keeping all raw condition-monitoring data local to its originating site, evaluated across simulated multi-site fault-diagnosis scenarios [26]. This federated-transfer combination is, in our view, close to the structurally correct answer for genuinely cross-platform, cross-organization fault prediction at scale, for essentially the same reason federated meta-learning has attracted attention in the broader IoT personalization literature: it lets a large, heterogeneous population of machines and operators collectively benefit from shared diagnostic knowledge without requiring any of them to give up control of their own raw operational data, which is very often the actual, non-technical obstacle to pooling fault data across an industrial platform in the first place.

7. Open Challenges

Several limitations recur across the literature synthesized here and deserve to be stated plainly. First, and most importantly, the overwhelming majority of the cross-domain transfer results reviewed above -- across nearly every architecture discussed, from DCTLN through the subdomain, adversarial, and meta-learning variants -- are evaluated on transfer tasks constructed from a small number of public university test-rig datasets, with "domains" typically defined as different load or speed settings on the *same* physical rig rather than genuinely different machines [10,12,14,19]. This is a reasonable and controlled way to isolate the effect of a domain-adaptation technique, but it systematically understates the difficulty of true cross-platform transfer, where target machines differ in manufacturer, bearing geometry, sensor mounting, and installation environment simultaneously, not just in one controlled operating variable. The handful of papers that do target named, specific industrial equipment classes directly -- wind-turbine gearboxes, hydrodynamic transmissions -- are a welcome and more realistic exception, but remain a small fraction of the literature relative to benchmark-rig studies [21,20].

Second, multi-source and genuinely heterogeneous-fleet transfer remains comparatively underdeveloped relative to single-source-to-single-target transfer. A real industrial platform does not offer one clean source domain to transfer from; it offers a messy population of machines at various ages, maintenance histories, and instrumentation qualities, some of which are more relevant to a given target machine than others, and the problem of automatically identifying and weighting the most useful sources -- rather than assuming a single, well-chosen source domain has already been specified by the researcher -- has received comparatively little dedicated attention within fault diagnosis specifically, even though the general multi-source domain adaptation literature provides some of the relevant statistical machinery [7]. Third, open-set recognition of genuinely novel fault modes, addressed directly by only a minority of the work surveyed here, deserves to be treated as a first-class requirement rather than an optional extension, because a platform-scale deployment will, with near certainty, eventually encounter failure modes absent from any source domain's labeled history, and a system that cannot flag

"unknown" is arguably more dangerous than no automated diagnosis at all, since it will confidently misclassify a genuinely new failure as a familiar, and probably less urgent, one [18].

Fourth, the interaction between transfer-learning accuracy and deployment-hardware constraints is, in our reading, still treated as two largely separate research conversations rather than one. The architectures that report the highest cross-domain diagnostic accuracy -- multi-scale convolutional networks, graph-convolutional adversarial models, CNN-Transformer hybrids -- are, almost without exception, evaluated without reference to parameter count, memory footprint, or inference latency on the class of edge hardware that an IIoT deployment would actually use, while the papers that do take edge deployment seriously tend to use comparatively simple classifiers (random forests, small distilled networks) that have not, to our knowledge, been evaluated against the strongest domain-adaptation techniques discussed in earlier sections on the same cross-domain transfer tasks [24,25,17]. Closing that gap -- benchmarking the field's best domain-adaptation architectures specifically under realistic edge memory and latency budgets, rather than as an afterthought -- would substantially strengthen the practical case for cross-platform transfer learning in fault diagnosis. Fifth, privacy-preserving, federated approaches to cross-organization diagnostic knowledge sharing remain represented by a small number of studies relative to how central a role data-governance concerns plausibly play in whether cross-platform diagnostic knowledge sharing happens at all in practice; the one detailed federated transfer study we found reports promising results, but a single positive result is not yet a mature subfield [26].

8. Conclusion

The literature reviewed here supports a fairly specific and, we think, well-evidenced claim: transfer learning and domain adaptation, applied to vibration-based and condition-monitoring fault diagnosis, reliably close a substantial part of the accuracy gap that opens up when a deep diagnostic model trained under one set of conditions is deployed under a genuinely different one, whether that difference is operating load and speed, physical machine, or equipment class [3,5]. The technical toolkit for doing this -- statistical distribution matching in the DAN and CORAL tradition, adversarial

domain-invariant representation learning, and finer-grained subdomain and class-aware alignment methods that correct the class-collapse failure mode of naive global alignment -- is by now mature and has been repeatedly validated on standard bearing and gearbox transfer benchmarks by independent research groups [6,7,12,14]. The field has also, more recently, begun pushing into the harder and more industrially relevant settings that a true cross-platform deployment requires: few-shot adaptation when only a handful of target labels exist, meta-learning approaches that treat "a new target machine" as one instance of a broader task distribution rather than a single fixed problem, open-set recognition of fault modes never seen during training, and federated approaches that let diagnostic knowledge travel across organizational boundaries without raw sensor data ever having to travel with it [19,18,26].

What keeps this pattern from being unambiguously deployment-ready today is less a shortage of algorithmic ideas than a shortage of evidence at realistic scale and under realistic constraints: most reported gains come from transfer between artificially defined domains on a small number of shared laboratory test rigs rather than from genuinely heterogeneous industrial fleets, multi-source transfer across many machines at once is comparatively underexplored relative to the clean single-source-to-single-target setting, and the architectures that achieve the strongest cross-domain accuracy are rarely evaluated against the memory and latency budgets of the edge hardware a real IIoT deployment would use [10,17,25]. None of these gaps argue against the underlying approach; if anything, the consistency with which independent groups, working on rolling bearings, gearboxes, wind turbines, and industrial transmissions alike, have converged on broadly similar transfer-learning architectures is itself evidence that the pattern is sound. What the field needs next is less a new algorithmic breakthrough than a harder, more deliberate push toward evaluation on genuinely heterogeneous, field-deployed equipment, under the communication, privacy, and compute budgets that real cross-platform industrial IoT deployments actually operate within.

9. References

1. Zhao, R., Yan, R., Chen, Z., Mao, K., Wang, P., Gao, R.X. (2019) Deep learning and its applications to machine health monitoring. *Mechanical Systems and Signal Processing*, 115: 213-237. DOI: [10.1016/j.ymssp.2018.05.050](https://doi.org/10.1016/j.ymssp.2018.05.050)
2. Lei, Y., Yang, B., Jiang, X., Jia, F., Li, N., Nandi, A.K. (2020) Applications of machine learning to machine fault diagnosis: A review and roadmap. *Mechanical Systems and Signal Processing*, 138: 106587. DOI: [10.1016/j.ymssp.2019.106587](https://doi.org/10.1016/j.ymssp.2019.106587)
3. Pan, S.J., Yang, Q. (2010) A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10): 1345-1359. DOI: [10.1109/TKDE.2009.191](https://doi.org/10.1109/TKDE.2009.191)
4. Zhang, A., Wang, H., Li, S., Cui, Y., Liu, Z., Yang, G., Hu, J. (2018) Transfer Learning with Deep Recurrent Neural Networks for Remaining Useful Life Estimation. *Applied Sciences*, 8(12): 2416. DOI: [10.3390/app8122416](https://doi.org/10.3390/app8122416)
5. Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W. (2010) A theory of learning from different domains. *Machine Learning*, 79(1-2): 151-175. DOI: [10.1007/s10994-009-5152-4](https://doi.org/10.1007/s10994-009-5152-4)
6. Sun, B., Feng, J., Saenko, K. (2016) Return of Frustratingly Easy Domain Adaptation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1). DOI: [10.1609/aaai.v30i1.10306](https://doi.org/10.1609/aaai.v30i1.10306)
7. Long, M., Cao, Y., Cao, Z., Wang, J., Jordan, M.I. (2019) Transferable Representation Learning with Deep Adaptation Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(12): 3071-3085. DOI: [10.1109/TPAMI.2018.2868685](https://doi.org/10.1109/TPAMI.2018.2868685)
8. Lu, W., Liang, B., Cheng, Y., Meng, D., Yang, J., Zhang, T. (2017) Deep Model Based Domain Adaptation for Fault Diagnosis. *IEEE Transactions on Industrial Electronics*, 64(3): 2296-2305. DOI: [10.1109/TIE.2016.2627020](https://doi.org/10.1109/TIE.2016.2627020)
9. Wen, L., Gao, L., Li, X. (2019) A New Deep Transfer Learning Based on Sparse Auto-Encoder for Fault Diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 49(1): 136-144. DOI: [10.1109/TSMC.2017.2754287](https://doi.org/10.1109/TSMC.2017.2754287)
10. Guo, L., Lei, Y., Xing, S., Yan, T., Li, N. (2019) Deep Convolutional Transfer Learning Network: A New Method for Intelligent Fault Diagnosis of Machines with Unlabeled Data. *IEEE Transactions on Industrial Electronics*, 66(9): 7316-7325. DOI: [10.1109/TIE.2018.2877090](https://doi.org/10.1109/TIE.2018.2877090)
11. Zhao, B., Zhang, X., Zhan, Z., Pang, S. (2020) Deep Multi-Scale Convolutional Transfer Learning Network. *Neurocomputing*, 407: 24-38. DOI: [10.1016/j.neucom.2020.04.073](https://doi.org/10.1016/j.neucom.2020.04.073)
12. Zhu, Y., Zhuang, F., Wang, J., Ke, G., Chen, J., Bian, J., Xiong, H., He, Q. (2021) Deep Subdomain Adaptation Network for Image Classification. *IEEE Transactions on Neural Networks and Learning Systems*, 32(4): 1713-1722. DOI: [10.1109/TNNLS.2020.2988928](https://doi.org/10.1109/TNNLS.2020.2988928)
13. Wang, B., Zhang, M., Xu, H., Wang, C., Yang, W. (2023) A

- cross-domain intelligent fault diagnosis method based on deep subdomain adaptation for few-shot fault diagnosis. *Applied Intelligence*, 53(20): 24474-24491. DOI: [10.1007/s10489-023-04749-4](https://doi.org/10.1007/s10489-023-04749-4)
14. Li, X., Jin, W., Xu, X., Yang, H. (2022) A Domain-Adversarial Multi-Graph Convolutional Network for Unsupervised Domain Adaptation Rolling Bearing Fault Diagnosis. *Symmetry*, 14(12): 2654. DOI: [10.3390/sym14122654](https://doi.org/10.3390/sym14122654)
 15. Javed, M., Ding, S., Yan, H., Ma, L., Kumar, T. (2026) Explainable Domain-Adaptive CNN-Transformer for Bidirectional Cross-Domain Bearing Fault Diagnosis. *Electronics*, 15(17): 3896. DOI: [10.3390/electronics15173896](https://doi.org/10.3390/electronics15173896)
 16. Wang, Z., Lv, Z., Chen, X., Zhang, H., Li, Z. (2026) Maximum Margin Local Domain Adaptation for Bearing Fault Diagnosis Under Multiple Operating Conditions. *Machines*, 14(4): 388. DOI: [10.3390/machines14040388](https://doi.org/10.3390/machines14040388)
 17. Jiao, X., Zhou, Y., Liu, X., Ma, Z., Xu, F. (2025) A Lightweight Fault Diagnosis with Domain Adaptation for Defected Bearings. *Applied Sciences*, 15(8): 4579. DOI: [10.3390/app15084579](https://doi.org/10.3390/app15084579)
 18. Lei, T., Pan, F., Hu, J., He, X., Li, B. (2025) A fault diagnosis method for rolling bearings in open-set domain adaptation with adversarial learning. *Scientific Reports*, 15. DOI: [10.1038/s41598-025-88353-1](https://doi.org/10.1038/s41598-025-88353-1)
 19. Zhang, Y., Han, D., Tian, J., Shi, P. (2023) Domain adaptation meta-learning network with discard-supplement module for few-shot cross-domain rotating machinery fault diagnosis. *Knowledge-Based Systems*, 268: 110484. DOI: [10.1016/j.knosys.2023.110484](https://doi.org/10.1016/j.knosys.2023.110484)
 20. Sun, D., Yang, X., Yang, H. (2025) A bearing fault diagnosis method for hydrodynamic transmissions integrating few-shot learning and transfer learning. *Scientific Reports*, 15: 19011. DOI: [10.1038/s41598-025-04543-x](https://doi.org/10.1038/s41598-025-04543-x)
 21. Guo, J., Wu, J., Zhang, S., Long, J., Chen, W., Cabrera, D., Li, C. (2020) Generative Transfer Learning for Intelligent Fault Diagnosis of the Wind Turbine Gearbox. *Sensors*, 20(5): 1361. DOI: [10.3390/s20051361](https://doi.org/10.3390/s20051361)
 22. An, J., Ai, P., Liu, D. (2020) Deep Domain Adaptation Model for Bearing Fault Diagnosis with Domain Alignment and Discriminative Feature Learning. *Shock and Vibration*, 2020: 4676701. DOI: [10.1155/2020/4676701](https://doi.org/10.1155/2020/4676701)
 23. Xu, J., Zhou, H., Fang, Y. (2022) Intelligent Roller Bearing Fault Diagnosis in Industrial Internet of Things. *Wireless Communications and Mobile Computing*, 2022: 1860946. DOI: [10.1155/2022/1860946](https://doi.org/10.1155/2022/1860946)
 24. Liu, P., Zhang, Y., Wu, H., Fu, T. (2020) Optimization of Edge-PLC-Based Fault Diagnosis With Random Forest in Industrial Internet of Things. *IEEE Internet of Things Journal*, 7(10): 9908-9918. DOI: [10.1109/JIOT.2020.2994200](https://doi.org/10.1109/JIOT.2020.2994200)

25. Wang, Y., Yu, Z., Wu, J., Wang, C., Zhou, Q., Hu, J. (2024) Adaptive Knowledge Distillation-Based Lightweight Intelligent Fault Diagnosis Framework in IoT Edge Computing. IEEE Internet of Things Journal, 11(13): 23156-23169. DOI: [10.1109/JIOT.2024.3387328](https://doi.org/10.1109/JIOT.2024.3387328)
26. Xue, X., Zhao, X., Zhang, Y., Ma, M., Bu, C., Peng, P. (2023) Federated transfer learning with consensus knowledge distillation for intelligent fault diagnosis under data privacy preserving. Measurement Science and Technology, 35(1): 015108. DOI: [10.1088/1361-6501/acf77d](https://doi.org/10.1088/1361-6501/acf77d)

10. Declarations

Funding

This literature synthesis did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Conflict of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported here.

Data Availability Statement

This article is a literature synthesis and generated no new experimental data; all findings discussed are drawn from the cited publications, which readers should consult directly for their own data and code availability statements.

Ethical Approval

Not applicable, as this work involved no human participants, animal subjects, or collection of primary experimental data.