

Intelligent IoT Systems for Civil Infrastructure Monitoring: A Research Roadmap

Elisa Bertino (✉ bertino@purdue.edu)

Purdue University West Lafayette

Mohammad Jahanshahi

Purdue University West Lafayette

Ankush Singla

Purdue University West Lafayette

Rih-Teng Wu

Purdue University West Lafayette

Research Article

Keywords: Autonomous IoT Devices, Deep Neural Networks, Analytics at the Edge

Posted Date: December 2nd, 2020

DOI: <https://doi.org/10.21203/rs.3.rs-119076/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published at Discover Internet of Things on February 24th, 2021. See the published version at <https://doi.org/10.1007/s43926-021-00009-4>.

Intelligent IoT Systems for Civil Infrastructure Monitoring: A Research Roadmap

E. Bertino · M. R. Jahanshahi
A. Singla · R.-T. Wu

Received: date / Accepted: date

Abstract This paper address the problem of efficient and effective data collection and analytics for applications such as civil infrastructure monitoring and emergency management. Such problem requires the development of techniques by which data acquisition devices, such as IoT devices, can: (a) perform local analysis of collected data; and (b) based on the results of such analysis, autonomously decide further data acquisition. The ability to perform local analysis is critical in order to reduce the transmission costs and latency as the results of an analysis are usually smaller in size than the original data. As an example, in case of strict real-time requirements, the analysis results can be transmitted in real-time, whereas the actual collected data can be uploaded later on. The ability to autonomously decide about further data acquisition enhances scalability and reduces the need of real-time human involvement in data acquisition processes, especially in contexts with critical real-time requirements. The paper focuses on deep neural networks and discusses techniques for supporting transfer learning and pruning, so to reduce the times for training the networks and the size of the networks for deployment at IoT devices. We also discusses approaches based on machine learning reinforcement techniques enhancing the autonomy of IoT devices.

E. Bertino
CS Department, Purdue University
E-mail: bertino@purdue.edu

M. R. Jahanshahi
Lyle School of Civil Engineering, Purdue University
E-mail: jahansha@purdue.edu

A. Singla
CS Department, Purdue University
E-mail: asingla@purdue.edu

R.-T. Wu
Lyle School of Civil Engineering, Purdue University
E-mail: jwu952@purdue.edu

Keywords Autonomous IoT Devices · Deep Neural Networks · Analytics at the Edge

1 Introduction

In future smart cities, many decision processes in critical infrastructure and emergency management will be based on machine learning (ML) techniques. One particular application is the processing of large datasets of visual images and other types of data for defect assessment where the data is collected by a swarm of IoT devices (devices, for short), some of which can be mobile. e.g., small unmanned aerial vehicles (UAVs), and robots. In this context, examples of defective regions are corrosion and cracks in buildings and facilities [3], and potholes on roads. A critical requirement for the success of such assessment processes is the reliable detection, quantification and localization of defective regions. Furthermore, in such applications, a real-time assessment is often critical so that the swarm can decide regarding the optimum strategy and corresponding actions for effective data collection in unknown environments, e.g., robots used for earthquake reconnaissance and rescue where they enter buildings whose plan is unknown to the robots. On the other hand, for such assessments to be reliable it is critical that data be of good quality since poor data may negatively affect the accuracy of classification and predictions, and consequently may introduce additional costs and time overhead.

In general, acquiring data and making sure that the data is of high quality, especially for real-time decisions, is expensive due to difficulties in reaching the regions where the objects of interest are located and the need for human-intensive assessment. However, today we have many technologies that can be leveraged to devise effective and inexpensive solutions, including: deep neural networks for image analysis; image processing techniques; devices able to acquire images and other types of data; 5G networks and edge computing processing [6]; crowd-sourcing [1]. Mobile devices such as drones are becoming quite powerful in terms of the sensing capabilities they offer – as an example the nano Black Hornet 3 drone is equipped with a microcamera core and a visible sensor to allow for enhanced image fidelity [11].

The use of devices for data acquisition, of course, is not new and almost all application domains we may think of use these devices, including civil infrastructures, smart cities, smart agriculture, emergency management, and environmental protection. However, the common practice is to use the devices just as a data collection means. Data is acquired by devices and then transmitted to some centralized large server, such as a cloud server, for processing and analysis. Such an approach may not always be optimal in many situations. Intermittent communications and communication disruptions, as in the case of battlefield and emergency management scenarios [7], may make it difficult to transmit data to a centralized server. In addition, in many such scenarios, it is critical to quickly analyze the data and, based on the analysis results, determine whether additional data needs to be collected or specific actions ex-

ecuted. For example, consider the case of the collapse of the Morandi Bridge in Genoa [5]. In such a scenario, being able to quickly detect anomalies in a bridge and notify incoming vehicles would save human life and a few seconds may make the difference. Approaches by which data has to be sent to a cloud server would not be viable.

Another important consideration is that for many applications it is critical to ensure that data be of good quality. For instance, in the case of image data, it is important that the image of an object of interest (e.g., a crack on a wall) would not be occluded. Completeness (e.g., making sure that no relevant data is missing) is equally critical. For example, if a failure is detected on one side of a structure, it is critical to determine whether the failure extends to the other sides of the structure. Current assessments regarding the quality and completeness of acquired data requires not only to send the data to remote servers, but often also to involve human analysts to evaluate the data and provide further data acquisition instructions to the remote devices. This approach is not effective and requires extensive real-time human involvement. Furthermore, this approach would not be scalable as the numbers of data acquisition devices increase. Additionally, as new phenomena of interest arise in the application domain of interest, it is critical to reduce the time required to deploy the needed data analytic solutions. For example, machine learning approaches that require training labeled data sets that are expensive and/or time-consuming to collect and label may not be effective for many scenarios. We thus need approaches by which:

- Devices can directly perform analyses on the collected data as when there are real time requirements, transmitting the analysis results requires significantly less communication bandwidth and, in addition, may allow the devices to quickly send high-priority safety information to humans, vehicles, and other parties.
- Devices can autonomously decide which data to collect based on data they have already collected and locally analyzed.
- The times required for generating analysis models deployed at the edge be minimized as much as possible.

In this paper, we discuss approaches addressing the above three requirements and outline a research roadmap. We focus on data analysis techniques based on deep convolutional neural networks (CNNs) and investigate the use of transfer learning (TL), so to minimize the time/cost for training the networks, and CNN pruning, so to reduce the size and inference costs of the networks, for deployment at devices. We discuss also approaches supporting autonomous and adaptive data collection based on the use of reinforcement learning (RL).

2 Analytics at Devices

In this section, we focus on the problem of efficiently training CNNs and deploying them at devices. We first summarize initial results from one of our




Fig. 1 Framework for device-based inspection system.

projects. We then briefly discuss related work and outline open research directions.

2.1 Initial Results

Figure 1 show the high-level steps of our approach for device-based failure detection [8]. Since typically multiple types of damage exist in infrastructures (e.g., cracks, corrosion, spalling, exposed rebars, etc. on a concrete surface), it is often very difficult and/or expensive to acquire and label sufficient data for network training. To address this limitation, we have adopted TL where a pre-trained deep CNN is used to detect a new type of damage. TL is a very popular choice for vision-based infrastructure assessment, since it requires less training data compared to when a CNN is trained from scratch. To this end, we have used very large networks that have shown success in the ImageNet Challenge. The drawback is that the number of damage classes for civil infrastructure assessment is far less than the number of classes in ImageNet (i.e., 1000 classes). This means that while the feasibility of TL has been acknowledged for health monitoring of civil infrastructures, this solution is not efficient (i.e., the networks are unnecessarily too deep for the problem of interest) and, consequently, TL is not suitable for detection at devices. To address this issue, we have used network pruning to enhance the resource efficiency for on-device analysis while still maintaining good detection accuracy. This approach allows one to deploy deep CNNs that are quite accurate, require low storage and computing resources, and can make decisions very quickly at devices.

We have tested this approach for the detection of crack [3] and corrosion [9] surface defects. The methodology starts with a pre-trained network (e.g., VGG16 [10]), and recursively reduces the network size by using the Taylor-expansion based pruning technique [12]. Since the pre-trained network is originally designed for the ImageNet 1000 image categories, it is very large in size and may contain redundant convolution kernels that do not contribute to the new detection problems of interest. The pruning technique evaluates the

importance of the convolution kernels and removes the kernels with the least contribution. After removing the kernels, the pruned network is fine-tuned again to enhance its performance for damage detection. Based on the detection performance, the user can determine whether or not to further prune the network following the same procedure.

In our experiments we started from VGG16, and used 29,468 crack, 29,780 non-crack, 33,039 corrosion, and 34,148 non-corrosion image patches to train and test the network. In this case, pruning stops if the detection accuracy after fine-tuning drops more than 3%. When up to 84% filters are pruned, the mean detection accuracy after fine-tuning for crack and corrosion is approximately 99% with a standard deviation below 0.5%. This demonstrates the robustness of the proposed approach as the variations in the performance are quite small when the pruned network still has the capacity to deal with the detection task. When 97% of the filters are removed (i.e., only 128 filters left), the mean accuracy of crack detection drops to 84.7% with a standard deviation of 19.86%, and the mean accuracy of corrosion detection drops to 96.0% with a standard deviation of 0.95%. This indicates that the pruning should be terminated due to the increasing variation and decreasing accuracy in detection performance. We also compared the inference time (i.e., the total time required to classify one 720×540 image) of VGG16 and ResNet18 [13] when deployed at devices (i.e., NVIDIA Jetson TX2 GPU) for damage detection. By removing 84% and 79% of the convolution kernels from VGG16 and ResNet18, respectively, the inference time for crack detection decreases from 279.7 (sec) to 31.6 (sec) for VGG16 and 36.8 (sec) to 8.9 (sec) for ResNet18. For the corrosion dataset, the inference time decreases from 275.7 (sec) to 30.6 (sec) and from 34.1 (sec) to 9.0 (sec) for VGG16 and ResNet18, respectively. In terms of memory reduction, the memory demands of VGG16 drops from 525 (MB) to 125 (MB), and the demands of ResNet18 drops from 44 (MB) to 2 (MB). By utilizing network pruning, VGG16 achieves a 89% reduction in inference time and 80% reduction in memory, while ResNet18 achieves a 76% reduction in inference time and 95% reduction in memory demands, without decreasing damage detection performance.

Our results indicate that network pruning is an important step towards incorporate deep learning architectures into devices. However, there are still several open questions that need further research. For instance, the selection of the appropriate pruning algorithm to reduce the network size. Also estimating the sensitivity of the pruning algorithm with respect to various network configurations is critical. The sensitivity should be considered in different aspects, e.g., inference performance and pruning efficiency.

2.2 Related Work

TL techniques have already been widely investigated [15], whereas pruning techniques have been receiving wide attention more recently. Notable related work includes:

- *Transfer Learning*: Zhu et al. [16] addressed heterogeneous TL and used information from text data to improve model’s performance in image classification. Aytar et al. [17] and Tommasi [18] addressed the deficit of training samples for some categories by adapting classifiers trained for other categories. Oquab et al. [19] showed that layers trained on ImageNet [20] can be reused to extract the mid-level features of images in the PASCAL VOC dataset. Shin et al. [2] addressed two specific computer-aided detection problems in medical images by fine-tuning CNNs pre-trained using a huge training set, such as CIFAR-10 [21], which has one million images from ten different classes. They further explored different popular CNN architectures and their performances on datasets of different sizes, concluding that the trade-off between learning more accurate models and using more training data should be carefully considered. Collobert et al. [22] explored TL for natural language processing. Hwangbo et al. [23] showed how to enhance RL by applying TL. More recently Singla et al. [14] developed a TL approach based on generative adversarial networks (GANs) for data different from images.
- *Model Pruning*: Network pruning has been investigated in the context of CNNs. The early work of LeCun et al. [27] proved that network pruning is a valid strategy to reduce the network complexity and over-fitting by using a diagonal Hessian-based approximation. In this kind of approximation, neurons were removed based on the result of calculations obtained from Hessian matrix. Recently, Han et al. [28] proposed an effective compression approach for CNNs. They have tested their approaches on both VGG16 and AlexNet [40] and on different hardware platforms. The experiments showed that their approach was able to substantially reduce the size of the networks without losing accuracy. He et al. [29] proposed a kernel pruning algorithm for CNNs. Their experimental results show that for VGG16 their approach was able to obtain a time speed up of 5X with only 03% error increase.

2.3 Research Directions

In many applications, such as civil infrastructure monitoring and emergency management, the amount of training data is limited and thus TL is essential to still obtain good models even with limited data. Pruning is then critical to reduce the size of the models obtained from transfer learning. However, combining these two techniques requires analyzing the optimal ordering of TL and pruning steps, and assessing the impact that different strategies would have on the performance of different CNNs. Equally important is to assess and optimize the computing resources required for inference at different edge devices. In what follows we discuss some relevant research directions.

2.3.1 Optimal ordering for the execution of pruning and TL steps

It is important to determine whether it is better to first execute TL and then pruning – the strategy adopted in our preliminary work – or vice-versa. Both these strategies could be beneficial. By applying TL first, one can make sure that models adapt well to the target dataset, and then pruning can remove redundant neurons or layers. In this approach, the probability of removing the right neurons or layers from the model is higher. However, if the size of the target training dataset is small, it is better to first reduce the size of the model by applying pruning. More neurons or layers in a model means more parameters, and thus the time for training and fine-tuning will increase as well as the inference times. In addition, training a large model with a small amount of data may cause model over-fitting. Consequently, if the target dataset is small, it is better to first remove redundant neurons or layers, and then train the network with the small dataset.

2.3.2 Optimal pruning strategy

There are three pruning strategies that can be adopted:

- *Neuron removal*. This is a common pruning strategy. There are various criteria for removing neurons:
 - Threshold-based pruning. This strategy analyzes the weight of neurons and removes those having a weight less than a threshold. Results by Han et al. [24] show that this strategy reduces the number of neurons by a factor of $9\times$ without incurring accuracy loss.
 - Taylor-based pruning. This strategy, used in our preliminary work, uses Taylor expansion [12] to determine the importance of neurons and thus allows one to remove the least important ones.
- *Layer removal*. Previous results by He et al. [25] have shown that most of the neurons in the middle layers of ResNet50 have zero weight. Such neurons not only do not perform any feature extraction, they may also result in information loss. Thus, it is possible to skip those redundant layers by the introduction of shortcut connections [26] that skip one or more layers.
- *Kernel removal*. Previous results by He et al. [29] have shown that kernel removal is another possible approach to prune CNNs, as an images have RGB three different colors that correspond to three kernels in the CNN architecture. Thus, in some applications not every kernel is needed. For example, if one needs detect red roses, the red kernel would play a more important role than other kernels. Therefore one can keep the red kernel and remove some redundant kernels. Most of the times, a CNN has more than RGB three kernels, so there are several redundant kernels that could be removed.

Experimental comparisons to determine which pruning strategy works better depending on specific datasets and networks are critical. In addition, an interesting approach to explore would be to combine the three different pruning

strategies: neuron, kernel, and layer removal. For instance, in ResNet50 after removing some layers, redundant/useless kernels and neurons could still be left in the network. Removing those kernels and neurons may further reduce the network size and inference time.

2.3.3 Multi-step Transfer Learning and Pruning

There can be scenarios where one might need multi-step TL and pruning (MSTLP). MSTLP refers to obtaining a target network, specialized for certain classification decisions, by performing multiple TL and pruning steps. In such an approach one would start from an initial general network. Then, an intermediate trained network is derived from the initial one by applying TL and pruning. From the intermediate network one can derive another intermediate network and continue this procedure. For example, one can use TL and network pruning on a network initially trained using a generic dataset, such as ImageNet, in order to obtain an optimized network model for civil infrastructure monitoring applications that can detect various defects such as cracks, corrosion, etc. One can then use this intermediate model to train specialized binary classification models that just detect specific categories such as corrosion/no-corrosion, crack/no-crack, etc. This multi-step approach has several advantages over a single-step TL approach that learns directly from a generic dataset: (a) TL and network pruning can be much faster if they start from an intermediate specialized network than if they start from a larger generic network trained for large number of categories not relevant to the application of interest. (b) The use of intermediate networks enhances flexibility in cases in which one needs to identify more than one class as well as use a specialized network when one just needs to detect one category of objects/defects. For instance, one might just be interested in only identifying cracks for civil infrastructure inspection purposes. (c) Detection accuracy is enhanced as MSTLP may help prevent overfitting. (d) The network size obtained from MSTLP will be smaller than the ones obtained from training specialized classifiers directly from the generic dataset. Investigating MSTLP is an interesting research direction.

3 Adaptive Data Acquisition

We now focus on the other crucial aspect of a smart health inspection system based on intelligent autonomous devices. Consider the case of an inspection device, such as a robot sent to inspect a civil infrastructure (e.g., a bridge). In this case, the device would carry some sensors (e.g., cameras) and perform on-board analyses using the approaches discussed in the previous section. The device will first need to identify a target of interest (e.g., crucial structural components of a bridge), navigate itself to a position closer to the target, and then collect data to determine the presence of damage as well as whether to proceed with a more detailed inspection. During the inspection process, the

device is constantly facing decision making problems whenever an input data sample is acquired. For instance, what should be the next movement of the device if no damage of interest is present in the input image? Is the device close enough to the target? Is there any damage detected on the target? Is there any other potential damage nearby the target that requires more data collection? Is the collected data of good quality enough for the device to make the next decision? In such a context where many decisions need to be taken, RL appears a suitable approach.




Fig. 2 The concept of RL.

RL, inspired by how human learning, is a technique to determine the optimal decision by interacting with the environment. Compared to the conventional Q-learning algorithm, which becomes extremely inefficient for large scale state-space problems [39], the incorporation of deep neural networks (DNNs) into RL makes the learning process able to deal with high dimensional problems. In what follows we refer to DNN-based RL as DRL.

3.1 Related Work

Recent advances in robotic navigation/control fields have demonstrated the great potential of DRL for smart inspection systems and other similar applications. Tai et al. [30] demonstrated the capability of using DRL to train a robot to reach a pre-defined target location without collision in a map-less environment. Cheng and Zhang [31] used DRL to navigate a boat to a target, and avoid obstacles on the way. Mirowski et al. [32] used DRL for training agents to navigate in large and visually rich environments with various starting points and destinations. Zhu et al. [33] developed a robotic navigation system for indoor scenarios using DRL. The robot is trained in a simulated environment and is able to navigate to user-defined indoor objects (e.g., a sofa or a desk) in the testing stage. Hwangbo et al. [23] used DRL to control a quadrotor to stabilize itself when subjected to extreme external forces. Moreover, DRL has been applied in various video games where a trained artificial intelligence (AI) system is able to outperform human players. Hasselt et

al. [34] proposed the double Q-learning algorithm to address the issue of over-estimation problem encountered by deep Q-network (DQN). The performance of double Q-learning was tested on multiple Atari games and it was shown to be effective. Wang et al. [38] proposed a dueling network architecture to decouple the estimation of value and advantage in DQN. The dueling network outperformed the double Q-learning algorithm [34] in the challenging Atari game. Schaul et al. [35] showed that by integrating experience replay into the double Q-learning algorithm proposed in [34], the performance of the AI system in playing Atari games can be further enhanced. In [36,37], DRL was used to play first-person shooter (FPS) games (e.g., Doom video game). The agent was trained to explore the map, collect items, search and fight against the enemies. Due to DRL's success in these challenging decision making problems, we consider DRL as the best fit for smart inspection systems.

3.2 An Example Inspection Scenario and a DRL-based Framework

Consider an inspection device that starts with an initial state S_0 and observes data D_0 (e.g., an image). Based on D_0 , the device selects an action A_0 from the action set A and moves to a new state S_1 . By interacting with the environment (or an environment simulator), an immediate reward R_1 is assigned to the device, and the device will move to new states by repeating the steps of data acquisition, selecting actions, and receiving rewards. The objective of RL is to find the optimal policy, i.e., the set of best actions based on observing the data, that maximizes the cumulative rewards for the achieve to achieve. During training, the underlying driving mechanism is the Bellman equation [39] given by the equation below, that updates the Q -value of picking an action a at state S_i :

$$Q(S_i, a) \leftarrow (1 - \alpha^k)Q(S_i, a) + \alpha^k [R(S_i, a, S_{i+1}) - \rho^k t(S_i, a, S_{i+1}) + \gamma \max_{b \in A(S_{i+1})} Q(S_{i+1}, b)]$$

In the previous equation, $R(S_i, a, S_{i+1})$ is the immediate reward after choosing action a at state S_i , $t(S_i, a, S_{i+1})$ is the transition time from state S_i to S_{i+1} , $\max_{b \in A(S_{i+1})} Q(S_{i+1}, b)$ is the maximum Q -value among all the actions $b \in A(S_{i+1})$ in the next state S_{i+1} , k is the iteration number, α^k is the learning rate at iteration k , ρ^k is the average reward at iteration k , and γ is the discount factor for the expected future rewards. At the early stage of training, the device has a higher tendency to choose the exploratory actions in order to discover the good actions in the state-action space. As the training proceeds, the device will gradually choose the greedy actions, i.e., the actions that have the highest Q -value. Note that the proposed DRL approach approximates $Q(S_i, a)$ with DNNs in order to account for the infinite number of state-action pairs in our problem. Figure 2 illustrates the fundamental concept of RL in the context of smart inspection systems.

In our inspection scenario, the device would first collect data at a farther distance to capture the overview of the building, and then would moved to a closer position to inspect the damage on a first floor column thoroughly. Such




Fig. 3 The example DRL framework.

scenario is reflected by our example DRL-based framework shown in Figure 3. A CNN is employed to extract features from the input image, and the features are sent to both the navigation and the damage detection network. The navigation network deals with the identification of the target of interest (e.g., first floor columns shown in Figure 4¹) and determines the movement of the device. The detection network identifies the presence of damage on the target and determine whether the data is of good quality (e.g., lack of blurriness and/or completeness), and whether to proceed with more data collection. Consider the image shown in Figure 4. The actions for the device to take can be defined as: $\{a_1 : \text{a column exists and move forward distance } d_1; a_2 : \text{a column exists and move forward distance } d_2 < d_1; a_3 : \text{a column does not exist and randomly picks the next movement direction and movement distance.}\}$ Since the device does see a column and should move much more to get closer to the column, the corresponding rewards can be assigned as: $\{R_1 = 1.0, R_2 = 0.5, R_3 = 0\}$ in order for the device to learn to move more when seeing images similar to the one in Figure 4. Note that the navigation and detection network will work jointly during the inspection, as the decision made by one network will affect the decision on the other. For instance, the navigation network will keep making decisions if the detection network cannot identify the presence of damage with Figure 4. Once the device moves to a closer position to the column, e.g., Figure 5, the actions for the detection network can be defined as: $\{a_1 : \text{a crack is detected and should slightly move to collect more data; } a_2 : \text{a crack is detected and requires no further data collection; } a_3 : \text{no damage is detected.}\}$ Since the device only captures a portion of the damaged column and should collect more data, the corresponding rewards can be assigned as: $\{R_1 = 1.0, R_2 = 0.5, R_3 = 0\}$ in order for the device to learn to collect more data when seeing images similar to the one in Figure 5.

Our DRL-based inspection framework has then to support two important tasks: 1) routine inspection - the device would be trained to perform regular inspections, in which the device scans through the structure and identifies all possibilities of damage presence, and 2) urgent inspection - the device would be trained for rapid assessment in the context of urgent inspection (i.e., inspections after a natural hazard) where the device will evaluate the severity

¹ The images in Figures 4 and 5 are snapshots of a video recorded by using a UAV by one of the authors of this paper. The UAV was operated to inspect a school building damaged during a strong earthquake which happened in southern Taiwan in February 2016.



Fig. 4 (The device is inspecting a building after earthquake.




Fig. 5 The device is inspecting the damage on a first floor column.

of damaged regions and it will perform more data collection around the damage of top priority (i.e, not a thorough inspection of whole structures).

3.3 Research Directions

3.3.1 DRL-based autonomous smart inspection system under two scenarios, i.e., routine inspection and urgent inspection.

Although DRL has been demonstrated to be effective in robotic navigation, there are challenges needed to be addressed in the context of health inspection. For instance, in the existing approaches the target/destination is predefined and will be identical in the testing stage, e.g., the items and the enemies in the Doom game [36]. In our context, the target of interest and the damage may have similar pattern but not identical as the training data in the testing stage. The sensitivity of DRL towards the changes in the objects of interest needs a thorough investigation.

3.3.2 Evaluation of the performance of different network configurations, i.e., DQN, double Q-learning and dueling network, on navigation and damage detection.

There are various network configurations proposed in previous work to enhance the performance of DRL in playing Atari games. However, there is no a particular configuration that outperforms the others in all categories of games [34]. Therefore, research is needed to determine which configurations are best for inspection systems.

3.3.3 Evaluation of the benefit of an environment simulator.

Unlike past work in which an environment simulator, such as a gaming engine or a predefined map, is available, our adaptive DRL data collection framework infers the information directly from the acquired training data. This is particularly useful for urgent inspection after natural hazards as the environment often changes drastically. An interesting research direction is to analyze whether an approximate simulator can bring additional benefits to the inspection system.

3.3.4 Transfer learning for reinforcement learning agents.

A critical challenge related to real-time requirements is that RL has high convergence times. To address this issue, an approach is to use TL techniques based on GANs [14].

The application of such a methodology to quickly bootstrap the DQN in the control RL system has two variations:

- *Reward Knowledge Transfer*: Under this variation, a few explorations are performed by the RL agent. Then a GAN is used to generate augmented or synthetic data (enhanced exploration data) by minimizing domain loss between the TL source domain and target domain. The resulting dataset

is then used to train the DNN of the RL agent at the target domain. A key environment specific parameter in this setting is the bias introduced while training the DNN of RL agent. For optimal convergence, the new exploration samples should be favoured over stale samples from the source domain. This approach requires two learning/training steps, one to train the GAN and another one to train the DNN of the RL agent at the target domain.

- *Quality Value Knowledge Transfer*: This variation is similar to previous one except that here we transfer knowledge about the RL model directly by exchanging Q values instead of exploration samples. Here the target dataset is generated from the baseline RL agent model (Q values) trained by some new explorations. Here the GAN has two tasks: minimizing domain loss and optimizing the Q function. The key goal here is balancing domain loss and quality loss while training the GAN. Since, there is only one learning step here, we need to introduce a bias (target domain over source domain) in the GAN itself; this can be done when selecting batches to train the GAN.

It will be interesting to compare the two approaches and investigate additional specific aspects related to TL for RL agents used for inspection systems.

4 Concluding Remarks

In this paper we have discussed the use of devices in a critical application domain, that is, the monitoring of civil infrastructures for defect identification and assessment. In the paper we have focused on the use of ML techniques to allow devices to carry out data analyses on-board and to enhance the autonomy of devices. Device autonomy is critical when dealing with emergency situation in which network communications can be become fragmented. The paper has discussed several research directions. However the area of ML for IoT systems is an important emerging application area with many other interesting research directions.

References

1. Mohammad Allahbakhsh, Boualem Benatallah, Aleksandar Ignjatovic, Hamid R. Motahari Nezhad, Elisa Bertino, Schahram Dustdar. Quality Control in Crowdsourcing Systems: Issues and Directions. *IEEE Internet Comput.* 17(2): 76-81 (2013)
2. Shin Hoo-Chang, Holger R Roth, Mingchen Gao, Le Lu, Ziyue Xu, Isabella Nogue, Jianhua Yao, Daniel Mollura, and Ronald M Summers. Deep convolutional neural networks for computer-aided detection. CNN architectures, dataset characteristics and transfer learning. *IEEE Transactions on Medical Imaging*, 35(5):1285, 2016.
3. Fu-Chen Che, and Mohammad R. Jahanshahi. Nb-cnn: Deep learning-based crack detection using convolutional neural network and naive Bayes data fusion. *IEEE Transactions on Industrial Electronics*, 65(5): 4392–4400, 2018.
4. Elisa Bertino, and Mohammad R. Jahanshahi. Adaptive and cost-effective collection of high-quality data for critical infrastructure and emergency management in smart cities - framework and challenges. *ACM Journal on Data and Information Quality (JDIQ)*, 10(1):1:1–1:6, 2018.

5. <https://en.wikipedia.org/wiki/PonteMorandi>.
6. Weisong Shi, Jie Cao, Quan Zhang, Youhuizi Li, and Lanyu Xu. Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5):637–646, 2016.
7. Elisa Bertino, Seraphin Calo, Maroun Touma, Dinesh Verma, Christopher Williams, and Brian Rivera. A cognitive policy framework for next-generation distributed federated systems: Concepts and re-search directions. *Proceedings of the 37th IEEE International Conference on Distributed Computing Systems, ICDCS2017, Atlanta, GA, USA, June 5-8, 2017*.
8. Rih-Teng Wu, Ankush Singla, Mohammad R. Jahanshahi, and Elisa Bertino. Pruning deep convolutional neural networks for efficient edge computing in condition assessment of infrastructures. *Comput. Aided Civ. Infrastructure Eng.* 34(9): 774-789 (2019).
9. Deegan J. Atha and Mohammad R. Jahanshahi. Evaluation of deep learning approaches based on convolutional neural networks for corrosion detection. *Structural Health Monitoring* 17(5): 1110–1128 (2018).
10. Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *Proceedings of the 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015*.
11. Flir launches next-generation black hornet 3 nano-uav; press release. <https://www.flir.com/news-center/press-releases/flir-launches-next-generation-black-hornet-3-nano-uav/>, 2018.
12. Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Kautz Jan. Pruning convolutional neural networks for resource efficient inference. *Proceedings of the 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017*.
13. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint, (arXiv:1512.03385)*, 2015.
14. Ankush Singla, Elisa Bertino, Dinesh Verma. Preparing Network Intrusion Detection Deep Learning Models with Minimal Data Using Adversarial Domain Adaptation. *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security, AsiaCCS2020, Taipei, Taiwan, October 5-9, 2020*.
15. Sinno Jialin Pan, Qiang Yang. A Survey on Transfer Learning. *IEEE Trans. Knowl. Data Eng.* 22(10): 1345-1359 (2010).
16. Yin Zhu, Yuqiang Chen, Zhongqi Lu, Sinno Jialin Pan, Gui-Rong Xue, Yong Yu, and Qiang Yang. Heterogeneous transfer learning for image classification. *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2011, San Francisco, California, USA, August 7-11, 2011*.
17. Yusuf Aytar and Andrew Zisserman. Tabula rasa: Model transfer for object category detection. *Proceedings of the IEEE International Conference on Computer Vision, ICCV 2011, Barcelona, Spain, November 6-13, 2011*.
18. Tatiana Tommasi, Francesco Orabona, and Barbara Caputo. Safety in numbers: Learning categories from few examples with multi model knowledge transfer. *Proceedings of the Twenty-Third IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2010, San Francisco, CA, USA, 13-18 June 2010*.
19. Maxime Oquab, Leon Bottou, Ivan Laptev, and Josef Sivic. Learning and transferring mid-level image representations using convolutional neural networks. *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*.
20. Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, Fei-Fei Li. ImageNet: A Large-Scale Hierarchical Image Database. *Proceedings of the 2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*.
21. Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (Canadian Institute for Advanced Research).
22. Ronan Collobert, Jason Weston, Leon Bottou, Michael Karlen, Korya Kavukcuoglu, and Pavel Kuksa. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12(Aug):2493-2537, 2011.
23. Jemin Hwangbo, Inkyu Sa, Roland Siegwart, and Marco Hutter. Control of a quadrotor with reinforcement learning. *IEEE Robotics Autom. Lett.* 2(4): 2096-2103 (2017).

24. Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. Proceedings of the Annual Conference on Neural Information Processing Systems, NIPS 2015, December 7-12, 2015, Montreal, Quebec, Canada.
25. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016
26. Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, Andrew Rabinovich. Going deeper with convolutions. Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015.
27. Yann LeCun, John S. Denker, and Sara A. Solla. Optimal brain damage. Proceedings of the Annual Neural Information Processing Systems, NIPS 1989, Denver, Colorado, USA, November 27-30, 1989.
28. Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. Proceedings of the 4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016.
29. Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. Proceedings of the IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017.
30. Lei Tai, Giuseppe Paolo, and Ming Liu. Virtual-to-real deep reinforcement learning: continuous control of mobile robots for mapless navigation. Proceedings of the 2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018.
31. Yin Cheng and Weidong Zhang. Concise deep reinforcement learning obstacle avoidance for under actuated unmanned marine vessels. *Neurocomputing*, 272(1):63 – 73, 2018.
32. Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andrew J. Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, Dharshan Kumaran, and Raia Hadsell. Learning to navigate in complex environments. Proceedings of the 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017.
33. Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J. Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. Proceedings of the IEEE International Conference on Robotics and Automation, ICRA 2017, Singapore, Singapore, May 29 - June 3, 2017.
34. Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.
35. Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. Proceedings of the 4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016.
36. Guillaume Lample and Devendra Singh Chaplot. Playing FPS games with deep reinforcement learning. Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.
37. Khan Adil, Feng Jiang, Shaohui Liu, Aleksei Grigorev, B.B. Gupta, and Seungmin Rho. Training an agent for FPS doom game using visual reinforcement learning and vizdoom. In *International Journal of Advanced Computer Science and Applications*, IJACSA 2017, Vol. 8, No. 12.
38. Ziyu Wang, Tom Schaul, Matteo Hessel, Hado van Hasselt, Marc Lanctot, and Nando de Freitas. Dueling network architectures for deep reinforcement learning. Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016.
39. Abhijit Gosavi. *Simulation-Based Optimization: Parametric Optimization Techniques and Reinforcement Learning*. Springer, New York, 2015.
40. Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. Proceedings of the 26th Annual Conference on Neural Information Processing Systems 2012, December 3-6, 2012, Lake Tahoe, Nevada, United States.

Figures




Figure 1

Framework for device-based inspection system.

a_e : exploratory actions
 a_g : greedy actions




Figure 2

The concept of RL.




Figure 3

The example DRL framework.



Figure 4

The device is inspecting a building after earthquake.




Figure 5

The device is inspecting the damage on a rusted column.