



# Predicting dynamic spectrum allocation: a review covering simulation, modelling, and prediction

Andrew C. Cullen<sup>1</sup> · Benjamin I. P. Rubinstein<sup>1</sup> · Sithamparanathan Kandeepan<sup>2</sup> · Barry Flower<sup>3</sup> · Philip H. W. Leong<sup>3</sup>

© The Author(s) 2023

## Abstract

The advent of the Internet of Things and 5G has further accelerated the growth in devices attempting to gain access to the wireless spectrum. A consequence of this has been the commensurate growth in spectrum conflict and congestion across the wireless spectrum, which has begun to impose a significant impost upon innovation in both the public and private sectors. One potential avenue for resolving these issues, and improving the efficiency of spectrum utilisation can be found in devices making intelligent decisions about their access to spectrum through Dynamic Spectrum Allocation. Changing to a system of Dynamic Spectrum Allocation would require the development of complex and sophisticated inference frameworks, that would be able to be deployed at a scale able to support significant numbers of devices. The development and deployment of these systems cannot exist in isolation, but rather would require the development of tools that can simulate, measure, and predict Spectral Occupancy. To support the development such tools, this work reviews not just the available prediction frameworks for networked systems with sparse sensing over large scale geospatial environments, but also holistically considers the myriad of technological approaches required to support Dynamic Spectrum Allocation.

**Keywords** Spectrum measurement · Dynamic spectrum allocation · Software defined radio · Machine learning

---

Benjamin I. P. Rubinstein, Sithamparanathan Kandeepan, Barry Flower and Philip H. W. Leong have contributed equally to this work.

---

✉ Andrew C. Cullen  
andrew.cullen@unimelb.edu.au

<sup>1</sup> School of Computer and Information Systems, University of Melbourne, Melbourne, VIC, Australia

<sup>2</sup> School of Engineering, RMIT University, Melbourne, VIC, Australia

<sup>3</sup> School of Electrical and Information Engineering, University of Sydney, Sydney, NSW, Australia

## 1 Introduction

Spectrum conflict is on the rise, due to the growth of devices seeking access to a finite amount of spectral resources. The competition for access to this spectrum already has a significant impact on innovation in both the public and private sectors. These costs will only grow as it is projected that there will be up to 50 billion internet enabled wireless devices by 2020 (Al-Fuqaha et al. 2015), attempting to exchange 50 exabytes of data per month by 2021 (Zhou et al. 2018). Supporting this growth will require a commensurate increase in capacity, through both capital investment and the development of new technologies. Failing to address this growth in demand will result in an increasingly congested spectral environment, cost increases, and access constraints that will affect mobile devices, the Internet of Things, satellites, and radar systems.

These issues can broadly be seen as a consequence of static spectrum allocation, which has produced a spectral environment that is simultaneously congested and underutilised (Shared Spectrum Company 2010). One approach for improving the efficiency of access to the spectrum environment would be to move to policies of Dynamic Spectrum Allocation, in which devices negotiate with the spectral environment for access, in a fashion that balances public and private needs against demand (Marčian et al. 2010). Building such systems around Software Defined Radios (SDRs) and Cognitive Radios, to allow devices to dynamically and intelligently manage their access to the spectral environment without necessarily requiring any centralised control (Iii 2000; Sithamparanathan and Giorgetti 2012), essentially performing local optimisation to improve the individual device experience (Jiang et al. 2017b).

Performing such optimisation is an inherently complex task, as it requires a deep understanding of the multifaceted nature of signal behaviour, and the interactions between signals and environment in which they exist. This behaviour is driven not just by technological factors—like the use of WiFi and cellular networks, or the level of data throughput expected by users—but also how we both move through the built and natural environments, and how the signals that we produce interact with said environments. Moreover, additional complexity results from how contributions to the signal environment can be driven by the state of the signal environment itself. While it is clearly impossible to predict the conscious and unconscious decisions that drive an individual's contribution to the spectral environment, it may well be possible to accurately predict the aggregate contributions of a population group.

Armed with the data to make such predictions and the appropriate methodologies to leverage said data, SDRs would significantly improve the efficiency of spectral utilisation. However, the pathway to deployment of such systems is particularly complex, due to the lack of representative data sets to test the accuracy of prospective predictive frameworks. While it would be possible to collect such data, the lead time to designing and deploying the systems required would be significant, which would in turn lead to significant costs. Furthermore, deploying fixed measurement systems prior to designing the predictive framework would mean that lessons learned from the predictive approaches could not be leveraged to improve the design of the systems used to measure Spectral Occupancy.

Clearly, a linear developmental approach, in which the systems engineering first considers measurement, then prediction, and finally the hurdles of deployment would require significant investments in both financial resources and time. Disrupting this sequential design process, and moving to a more efficient, circular design methodology, is possible by introducing simulated data to the design process. By constructing a simulacra of the spectral

environment, the performance of predictive techniques, the requisite engineering specifications for the sensors, and the designs of the sensor networks themselves can all be considered in parallel. While constructing such a simulacra is not without its challenges—due to the multifaceted nature of spectral occupancy, and the geographic and temporal scales over which it varies—doing so has the potential to significantly decrease the development time required to build the systems required to support Dynamic Spectrum Allocation.

### 1.1 Summary of contributions

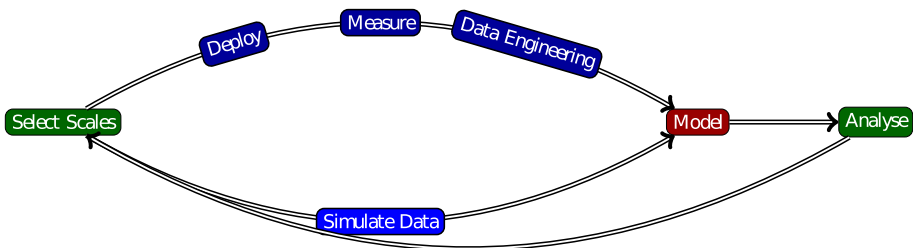
The task of predicting dynamic spectrum allocation intuitively feels as if it would be a natural extension of predictive modelling in other fields. However, the spectral and spatio-temporal dimensionality, the scale of underlying processes, and the inherent difficulties imposed by the dimensionality of the underlying data present unique challenges that do not have many analogues in related fields. Even when considered in the context of these sub-problems, the strong inter-dependencies between problem facets exacerbates the overall complexity. The nature of these inter-dependencies—as is emphasised within Fig. 1— inherently means that a good solution to one-sub-problem may be infeasible in the context of the broader problem space.

As such, this review considers a holistic approach the systems and techniques required to support dynamic spectrum allocation. This includes discussing the motivation of the problem space in Sect. 2; the measurement of spectrum usage in Sect. 3; constructing future predictions within Sect. 4; and the deployment of such systems in Sect. 5.

In taking such a holistic perspective, the contribution of this work is not only to review the performance of techniques (drawn from works within and outside the field of dynamic spectrum allocation), but also to provide a big-picture perspective to researchers considering facets of the spectrum allocation problem. It is hoped that by doing so future work can be directed in a fashion that is complimentary to the wider problem space.

In aide of this, our work makes the following specific contributions:

1. A concise definition of Spectral Occupancy, that can be applied to a range of domains, while being well suited for spectral–spatio–temporal predictive modelling, as covered within Sect. 2. This chapter also reviews the considerations and techniques required for the measurement of Spectral Occupancy.



**Fig. 1** Flowchart emphasising the inherent inter-dependencies in attempting to support Dynamic Spectrum Allocation. Here fundamental problem scales are selected, prior to either simulating or measuring data. This data is then modelled, and based upon the outcomes of this process, the fundamental assumptions regarding the appropriate scales may need to be revised

2. A comprehensive review of how traditional and machine learning techniques can be applied to predicting Spectral Occupancy. This can be found within Sect. 4, which outlines a novel taxonomy of prediction frameworks. This work is the first to consider holistically consider all elements of spectral–spatio–temporal prediction, including how these predictions are dependent upon sensing systems.
3. Due to the cost of collecting appropriate data for testing prediction regimes, this work provides the first review of techniques available for constructing synthetic data for spectrum related tasks. Section 5 outlines how this data can be generated, and how it can be utilised to significantly reduce the risk in developing and deploying large scale Occupancy predicting systems.
4. Also contained within Sect. 5 is a review of the engineering requirements for building, deploying, and managing systems to support SDRs at scale.
5. Finally Sect. 6 outlines both open questions and directions of research priority to support the task of Dynamic Spectrum Allocation.

## 2 Motivation and background

To date, studies of dynamic spectral allocation have typically focused upon specific sub-problems. One common feature is the reliance upon univariate signals, representing the capture of spectral data at a single point of space, over some defined time scale. These issues persist across both works considering the measurement of Spectral Occupancy (Xue et al. 2013; Subramaniam et al. 2015) and those testing the applicability of predictive regimes (Wang and Salous 2011).

While such work does demonstrate the applicability of localised spectral sensing, it inherently misses opportunities to understand key drivers of the signals environment, especially in urban environments. In such spaces, there is both the potential for there to be inherent correlations between sensor locations, and an inherent interdependency between the use of spectrum and the surrounding built environment. As such, rather than consider spectrum allocation at a single point in space, this work is the first to consider not just the suite of potential approaches, but how these approaches can be built into a cohesive, holistic predictive system.

In aide of this, we begin by defining how Spectral Occupancy should be considered quantitatively. Any measure of Spectral Occupancy should reflect the difficulty of accessing the spectral environment. This environment is inherently inhomogeneous across not just time and space, but also the frequency at which a device is able to communicate. To support Dynamic Spectrum Allocation, the metric should also reflect not just the level of congestion within the spectral environment, but also the ability for the device to transmit signals relative to such congestion. In doing so, the measure should reflect that different devices will have a different perspective on what constitutes occupancy, as a mobile phone has far less power, and thus ability to access and impact congested spectrum, than a television transmission tower.

The Spectrum Occupancy  $O$  at any point of space can be considered as a time series that represents the proportion of the time window  $(t, t + \delta t)$  that a signal would face congestion from another transmitter, whose received signal would be above the considered threshold power in the bandwidth region of interest. This corresponds to the probability that *measured* signals at a point and time will be strong enough to induce congestion within the spectrum. Such a measure is consistent with the latest 802.22 wireless standard for

SDRs (Cordeiro et al. 2005). Building upon the Duty Cycle model of Spectral Occupancy (Wellens and Mähönen 2010) to incorporate spectral and spatiotemporal variation, we propose measuring occupancy as

$$O(t, \mathbf{x}, b, \hat{P}) = \frac{1}{\delta t} \int_t^{t+\delta t} \mathbb{1}_{\hat{P}}(P(\tau, \mathbf{x}; b)) d\tau \text{ where} \tag{1}$$

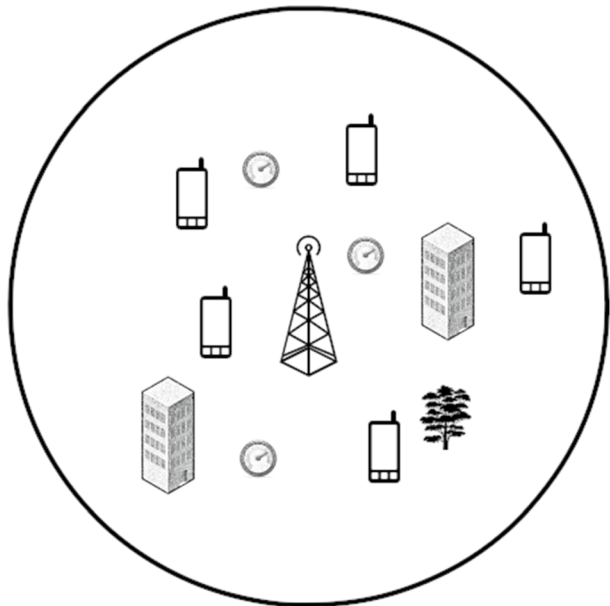
$$\mathbb{1}_{\hat{P}}(P(\tau, \mathbf{x}; b)) = \begin{cases} 1 & \text{if } P(\tau, \mathbf{x}; b) > \hat{P} \\ 0 & \text{if } P(\tau, \mathbf{x}; b) \leq \hat{P}. \end{cases}$$

here  $t$  and  $\delta t$  are the time and the length of the time window respectively,  $\mathbf{x}$  is the location at which the occupancy is to be measured,  $P$  is the maximum received power at a point in space and time,  $\hat{P}$  is the threshold power (in dBW), and  $b$  represents the bandwidth representing the frequency span  $[f_0, f_1]$ . Treating this as an integrable measure only with respect to time, and not space, is a consequence of imposing that the measure of occupancy relates to received power at any point. Moreover, considering the occupancy in terms of bandwidth, rather than frequency, reflects the potential for frequency modulation to occur, and also allows a degree of quantisation to be introduced to the model, which can reduce the complexity of the task of predicting Spectral Occupancy while simultaneously reflecting the nature of different device behaviours.

### 3 Measuring spectral occupancy: spectral sensing

Before considering the prediction of spectral occupancy, it is important to understand how deployed systems are able to capture the spectral environment. An exemplar sensing environment is shown in Fig. 2, in which a sparse network of sensors are able to capture the local spectral occupancy. These can be considered as local observations at

**Fig. 2** An example sensing domain, containing a base station antenna, four mobile devices, three fixed sensors, and both natural and constructed environmental factors



distinct  $\mathbf{x} \subset \hat{\Omega}$  at time  $t$ . Depending upon the context of the deployed system, any predictions may need to be constructed at either the discrete sensor locations  $\mathbf{x} \subset \hat{\Omega}$ ; or across a broader domain  $\mathbf{x} \subset \Omega$ .

To support the development of these systems, it is important to understand the behaviour of sensors of Spectral Occupancy, as their physical behaviours will inform the range of available prediction methodologies. The tools for such spectrum sensing fall into two broad categories, known as narrowband or wideband sensors (Sitham-paranathan and Giorgetti 2012), with the demarcation between the two determined by the number of frequency channels the sensor can access. Narrowband sensors consider one distinct frequency at each point in time, whereas wideband sensors aim to analyse a frequency bandwidth, in a manner that operates similarly to a number of narrowband sensors operating in parallel (Ali and Hamouda 2016), which introduces complexity to the implementation, but requires less energy while producing faster results than sequential-sensing (Lu et al. 2017).

Narrowband sensors can be constructed based upon a number of frameworks, with matched-filter detection (Veen and Van Der Wiellen 2003), energy detection (Sun et al. 2010), and cyclostationarity feature detection (Sutton et al. 2008) being the most commonly implemented. Other narrowband sensing techniques include covariance based detection (Zeng and Liang 2008) and cooperative machine learning based sensing (Wang and Yang 2016). With Matched-filtering, the received signal-to-noise ratio is maximised under the assumption that the primary user signal can be demodulated, which imposes additional constraints. Energy detection is a simplification of Matched-filtering, that introduces an energy metric to perform sub-optimal non-coherent detection in a manner that can require fewer samples to meet a detection constraint as compared to matched-filtering. Cyclostationarity feature detection is constructed around the assumption that detected signals contain periodic components, that can be leveraged to decrease the number of samples required for a detection event. However, the process of Cyclostationarity feature detection is computationally expensive, due to the need to both make multiple frequency transforms and calculate the signal's autocorrelation.

To support Dynamic Spectrum Allocation, sensing systems require the ability to provide information about wide swathes of the frequency range, in order to exploit extant opportunities. While techniques have been proposed for capturing a representative sample of the environment through a distributed network of sensors (Blum et al. 1997), they are still fundamentally limited in the amount of information that they can capture. In contrast, wideband sensors are able to sample the broader spectral environment with a self contained device, producing consistent and coherent results, without any of the potential robustness concerns stemming from a reliance on data from a series of networked narrowband sensors.

While the ability to capture larger swathes of the spectral environment with a single sensor is inherently appealing for SDRs, conventional wideband spectrum sensing often leads to a sampling rate so high that it becomes unaffordable, or introduces implementation issues that significantly increase the complexity of designing and integrating these systems. As such, novel signal processing approaches have been required to leverage the potential of wideband sensing. As channels in the wideband frequency spectrum are inherently sparse (Do et al. 2004), there is a clear application case for sparse sampling and Compressed Sensing (CS) (Candes et al. 2006). While most frequency-sampling approaches assume that the ability to reconstruct a signal is limited by the Nyquist frequency, CS incorporates knowledge of the signals sparsity to reconstruct signals

sampled significantly below the Nyquist frequency. This is made possible by formulating the frequency reconstruction problem as an under-determined linear system, and then using  $L^1$ -norm minimisation to construct a single solution to the under-determined system.

Within a signals sensing framework, sampling below the Nyquist frequency means that significantly fewer samples are required to accurately reconstruct a signal operating at a given frequency, which has implications for both the speed of processing, the temporal resolutions that are possible, and the hardware requirement for spectrum sensing (Sharma et al. 2016). Because of this, numerous authors have considered applications of CS to narrowband and wideband sensing (Tian and Giannakis 2007; Ma et al. 2016). Each of these authors have presented variants of sub-Nyquist wideband sensing that leverages CS, and while they each have their own relative strengths, a particularly promising approach can be found in one-bit sensing (Laska et al. 2011), which considers only the sign information of measurements, in a manner that significantly reduces the amount of data that needs to be stored and processed (Jacques et al. 2013). By treating these one-bit sign measurements as constraints for the CS based reconstruction, sparse signals can be reconstructed with high probability (Li et al. 2018). By discarding unnecessary components of the original signal, one-bit compressive sensing is able to perform fast sampling without growth in the computational and hardware complexity, while still being robust to noise.

Alternate approaches to sub-Nyquist sampling include wavelet detection, in which a continuous signal is decomposed into a series of markers, corresponding to changes in occupancy of the frequency band (El-Khamy et al. 2013). Other options include multi-band joint detection (Quan et al. 2008), adaptive thresholding (Gorcin et al. 2010) and filter bank detection (Kim and Takada 2009), although each of these techniques is limited by the necessary sampling rate, and the associated latency, complexity, and energy consumption that comes with sampling at the Nyquist frequency level.

Ultimately, the choice of narrowband or wideband, and the specific sampling technique within those frameworks, is a multifaceted optimisation task subject to competing objectives. For a system designed to be integrated into a broader machine learning modelling and forecasting framework, important factors will include the robustness to false positives, as well as both the sensitivity to noise and the robustness of the sensor to noise. Similarly, computational requirements, problem specific latency, and resource utilisation considerations must also be taken into account.

Another important consideration in the specification of sensor networks is the trade off between accuracy and efficiency. In these terms, accuracy relates to both the ability of a sensor to detect transmissions above its noise floor and separate out signal detections from noise, and the false positive rate of transmissions; while efficiency refers to the proportion of time the sensor can query a part of the spectrum, given both observation and throughput constraints (Yu et al. 2012). These two factors of the sensors performance are intrinsically coupled—as more focus is placed upon the accuracy of the sensor, system resources are consumed that reduce its throughput, which will in turn affect the efficiency of the sensor. However, it has been shown that by acting in parallel, networks of sensors can improve both accuracy and efficiency by sharing some of the sensing workload (Xie et al. 2010).

### 3.1 Distributed sensor networks

A network of sensors offer more opportunities than just improving the local performance of sensors—distributed sensor networks are a crucial tool for sampling geographic regions



that would be beyond the scope of an individual sensors capabilities. The design of such networks is inherently complicated, due to both the domain complexity—in terms of both the spectral, natural and built environments—and the range of objectives that must be considered. These objectives for resource allocation can include, but are not limited to, considerations of coverage, energy efficiency, communication and processing constraints, detection reliability, and robustness to failure modes, as well as more application specific concerns.

As the domain of interest increases in size, resource constraints make perfect sampling of the domain to be an almost impossible task. Instead, it is likely that there will be geographically correlated problem areas in the domain coverage, giving rise to regions known as holes (Ahmed et al. 2005). A coverage hole is a geographic region that is not covered by at least  $k$  sensors, where  $k$  is the specified level of coverage required from the application, to capture all data and to account for fault tolerance (Huang and Tseng 2005). Constructing accurate measurements of the geographic distribution of spectrum occupancy—including potential geolocation of signal sources—may require  $k \geq 3$  (Niculescu and Nath 2003). Peer-to-peer network topologies may also require multiple coverage to ensure robust connectivity. The design of the network may also lead to the formation of routing holes, in which network messages from a geographic region are unable to be communicated to the broader network, due to bandwidth constraints, missing links, or an improperly designed adaptive routing framework (Akkaya and Younis 2005).

Static distributed sensor networks designs can be broadly categorised as belonging to either Target, Area, or Cell-Edge (or Barrier) coverage. For Target coverage, the sensed domain has predetermined discrete target locations that have been marked as being particularly important; Area coverage attempts to maximise the sampling across a defined geographic region; and Cell-Edge coverage attempts to sense transitions into and out of a region of interest, without focusing upon the interior region. Of these three taxonomies, Area coverage has the greatest relevance to the problem of spectral occupancy prediction, and as such we shall focus our consideration of network design on techniques suitable for Area coverage.

Under a static regime, techniques from computational geometry like Voronoi cells (Vieira et al. 2003) and Delaunay triangulation (Wu et al. 2007) were initially popular for managing grid distribution, however they are limited in their ability to be adapted to manage multifaceted objectives. To maximise the coverage of a sensor network, numerous iterative schemes have been successfully tested, including Ant colony optimisation (Liu and He 2014), Glow-worm swarm optimisation (Liao et al. 2011) and Genetic Algorithms (Yoon and Kim 2013), with similar techniques also being employed to maximise the connectivity between network nodes (Younis and Akkaya 2008). Connectivity maximisation is particularly important in the design of systems that are robust to the impact of environmental or adversarial interference.

The design objectives of physical implementations of wireless sensor networks are more likely to be multi-faceted, and as such approaches that can consider balanced objectives must be considered as being of particular importance. In doing so, coverage and connectivity considerations can also be balanced with other design parameters, like energy efficiency, redundancy, and the expected lifespan of the network. Due to their well established efficacy for multi-objective functions, Evolutionary and Genetic algorithms (Khalesian and Delavar 2016), Particle Swarm optimisation (Pradhan and Panda 2012) and even predator scent marking and bee-colony algorithms (Hashim et al. 2016) have all been proposed as mechanisms for managing the deployment of static wireless sensor networks. The motivation for employing these techniques is their well known ability to handle multi-objective



optimisation tasks, however they do introduce computational burdens, and are inscrutable when attempting to determine the factors that have given rise to the suggested topology.

### 3.2 Mobile distributed sensing

Distributed sensing is not necessarily limited to fixed topologies—such sensors could be mounted on moving human agents, ground or aerial vehicles, or satellites. As such, (Heo and Varshney 2003) introduced both algorithms, and performance metrics for ensuring that self-deploying wireless sensor networks best meet objectives for coverage, uniformity, deployment time, and distance travelled. However, accessibility for perfect placement is not necessarily guaranteed, and as such deployment algorithms have been designed to account for uncertainty derived from aerial (Corke et al. 2004) or ground based (Chang et al. 2008) deployment systems. The performance of such a system can be further enhanced if the behaviour of individual sensors can be modified, either in terms of their sleep-wake cycles and sensor power (Abo-Zahhad et al. 2016), or using mobility to rearrange the sensor nodes (Niewiadomska-Szynkiewicz et al. 2016). Beyond relocation, node mobility can also be exploited to enhance coverage (Liang and Ren 2005), and to generate efficient data transfer structures (Zhan et al. 2017). The balance of these factors will then be driven by the resources available for deployment.

While observations of the spectral environment are critically important, their value is not able to be capitalised upon without having the supporting hardware and software to manage the data, and to support the post-processing required to generate predictions of future spectral utilisation. This task is inherently complex, due to the potential presence of structural heterogeneities in the collected data, differential delivery rates due to system latencies, and the need to be robust against adversarial attacks and security concerns (Lu et al. 2013). The implementation hurdles associated with these concerns grow nonlinearly with the size of the sampled region, as a consequence of the volume of data that must be processed and managed.

### 3.3 Robust network design

While the aforementioned holes can broadly be considered failures in design and implementation, there is also the possibility that they were introduced through the actions of outside actors, behaving in an inadvertent or adversarial manner (Karlof and Wagner 2003). By compromising a node in such a manner that its behaviour is disrupted, adversarial agents can generate denial-of-service conditions, or inject false data into the detection framework. A sink hole is one in which a node that has either been compromised or exploited introduces behaviour that increases the resource demands placed on neighbouring nodes, in terms of energy consumption or bandwidth. Such an attack aims to introduce a failure mode that propagates through the nodes, to induce an artificial coverage or routing hole (Kalita and Kar 2009).

Another classical attack would be jamming the communication between nodes (in a deliberate or inadvertent manner), to ensure that any information describing a detection cannot be propagated out from the jammers domain of influence. As a consequence of this, any detections that would be able to be identified by network inference would also be denied. Such an event is known as a black hole, and there are also gray holes, in which the malicious node intermittently behaves as a black hole. An alternate adversarial approach is the insertion of a wormhole (Pathan et al. 2006), in which malicious nodes distributed

across a wide geographic region create network tunnels between each other, with the aim of introducing erroneous routing behaviour, and, potentially, spurious geographic cross-correlates that could present as false detections, or hide true detections, affecting the accuracy of the sensors.

## 4 Predicting spectral occupancy

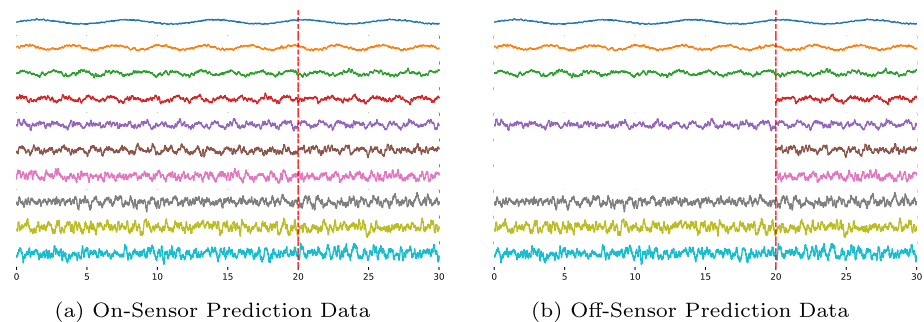
The complexity of predicting spectral occupancy (in the form of Eq. 1) is complicated due to the high-dimensional spectral–spatio–temporal nature of occupancy. This complexity only grows as the evolution of the spectral environment is dependent upon a large number of confounding factors, including the environment that the signals exist in—that can include built and natural topography, as well as weather events—and human factors. As such, it is important that consideration is given to prediction frameworks that can accurately incorporate sources of data beyond observations of the historic spectral environment.

Another factor that underscores the importance of reviewing available techniques is the local nature of occupancy observations, and the likely need in an SDR context for global (over a domain of interest) predictions. Given the cost of performing observations of the spectral environment it is likely that the observations are highly sparse—at a level beyond what has been required for other problem domains—relative to the resolution required for predictions.

In the context of these factors, which introduce challenges unique to this problem domain, the review of appropriate prediction techniques will be split into a taxonomy that independently considers predictions at, and away from observation sites as on- and off-sensor prediction respectively. This taxonomy can be considered as a generalisation of the consideration of prediction domains in Sect. 3. From a time series perspective, a basic exemplar of the differences between on- and off-sensor predictions can be seen in Fig. 3.

For on-sensor predictions, the historic values of a single sensor, or a network to predict the future state of spectral occupancy. By decomposing the task further to consider discrete frequency bands, the prediction task can be expressed as an  $m \times n$ -dimensional time series, where  $m$  and  $n$  are the number of connected sensors and frequency bands respectively.

As frequency bands can be treated as being nominally independent, the prediction task can be further decomposed to being  $n$  distinct  $m$ -dimensional time series. While the overall



**Fig. 3** Example of the information available for on-sensor and off-sensor predictions. Here the 10 time series provide historical information up to the red dotted line, and the prediction methodologies need to predict all subsequent points in the time series

number of time series is consistent across either of these approaches, the latter should be more computationally efficient, as the search for correlations within the time series should be proportional to the dimensionality of the time series. Unfortunately many suitable prediction frameworks are only designed for 1-dimensional time series, requiring the problem to be further decomposed, and to treat each sensor and frequency band as producing its own, distinct time series. While taking such an approach does prevent the prediction framework from leveraging information from neighbouring sensors or correlated frequency bands, it does significantly increase the range of approaches that can be applied to the task of predicting the spectral occupancy. While this lost information may be significant in some sensor network designs, as the spacing between sensors is increased the correlations will likely decrease, lending support for this decomposition approach.

On-sensor predictions can be considered as predictions of the form

$$O(t, \mathbf{x}_i, b_j, \hat{P}) = f_{i,j}(t) \quad (2)$$

for  $i \in m$  and  $j \in n$ , where each  $f_{i,j}(t)$  can be considered independent of every other sequence. It must be noted here that while this framework implies that each sensor is spatially stationary at  $\mathbf{x}_i$  across time, it is also possible to consider moving sensors under this framework. However, as many of these techniques are unable to incorporate contextual information—like position—into their temporal predictions, it must be expected that prediction quality will likely significantly decrease as the speed of motion increases, and becomes less deterministic and patterned.

Off-sensor predictions can be considered either as a distinct task, or a sequentially dependent extension of on-sensor predictions, in which the on-sensor extensions are extended across the spatial domain. The simplest approach for such extensions would be to use interpolation approaches to construct a geospatial representation of spectral occupancy across a domain of spatial interest.

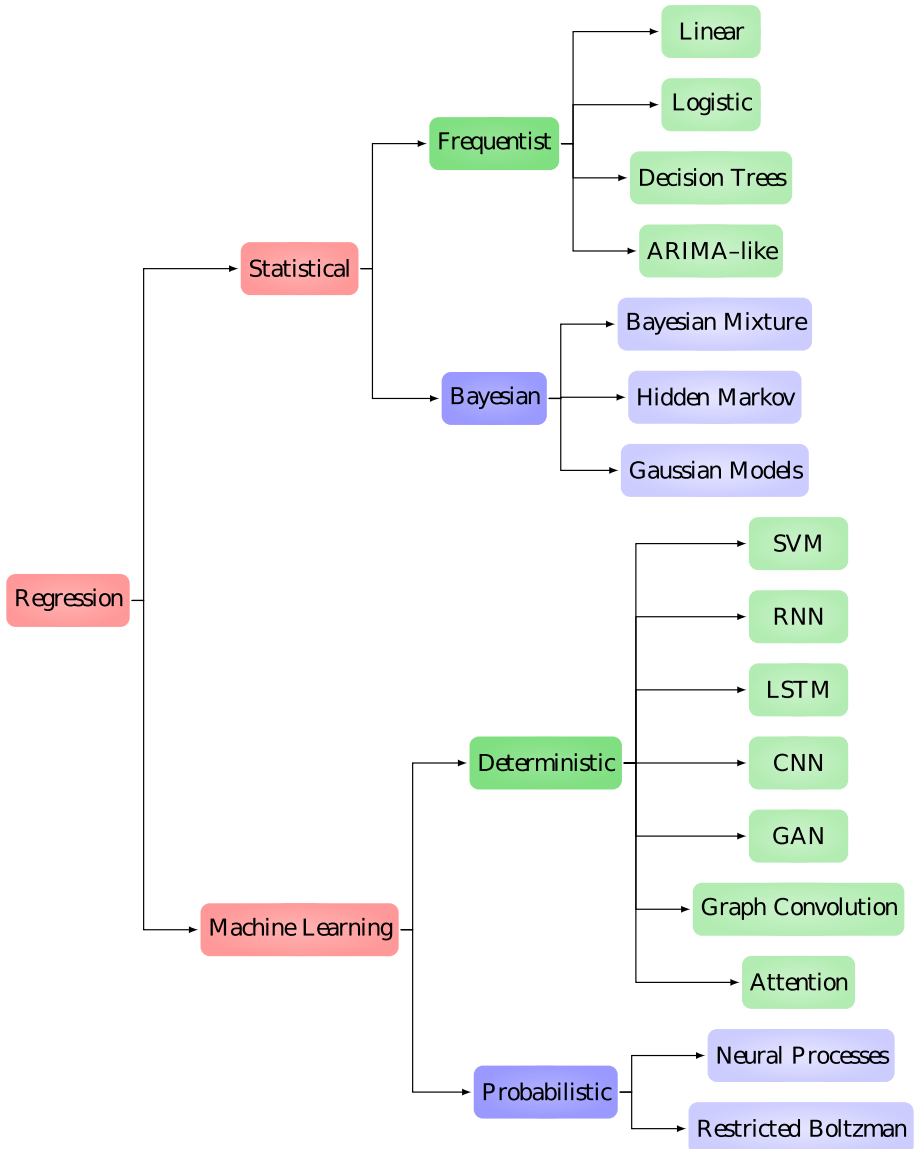
A complicating factor for both on- and off-sensor predictions is the behaviour of the sensors involved. As was discussed in Sect. 3, some sensor designs—particularly wideband sensors—are only capturing a subset of the frequency domain at any one time, as they need to perform a sweep of frequency space to capture the totality of spectral data. This results in a time series sequence that is inherently sparse. Depending upon the properties of the sensor, it may be possible to construct a continuous sequence by down-sampling from the sensor frequency. However, if this is not possible, then even on-sensor predictions may need to be considered using techniques more suitable for off-sensor problems, due to the inherent need for data imputation.

## 4.1 On-sensor predictions

Forecasting Spectral Occupancy can take numerous forms, depending upon both the nature of the requirements, and the resources available to support the prediction. In applications where the measurement of the spectral occupancy occurs at a single sensor—or where the sensors are independent and do not communicate with each other—the prediction task is well suited to the wide range of one-dimensional time series forecasting frameworks. Analogous problems have been successfully considered through statistical methods, ranging from linear extrapolation and sequential models like exponential smoothing (Gardner Jr 2006), through to more complex frameworks like ARIMA and Hidden Markov Models. Machine Learning offers another potential avenue for considering these one-dimensional

forecasting problems, with Recurrent Neural Networks and Long-Short Term Memory Networks showing particular promise.

The performance of these systems can be improved for systems of multiple sensors, if the sensors are able to pass information between one another. In this case, the system becomes analogous to a multidimensional time series problem, of a form that is well suited for machine learning prediction tasks. The remainder of this subsection will explore suitable techniques for single- and multi-dimensional time series predictions, following the



**Fig. 4** Taxonomy of regression algorithms for on-sensor wireless spectrum occupancy prediction

taxonomy of Fig. 4, with particular emphasis upon techniques that have been applied to analogous problem spaces.

#### 4.1.1 Statistical methods for temporal prediction

Historically, Bayesian models have frequently been used for predicting time series problems from varying domains, including spectral processes (Eltom et al. 2018a). The depth of a well-fit Bayesian explanatory model allows for explicit claims about the generative process that actually created the observation. Hence, the estimated latent generative model can synthesise observations based on the conditional probability relationship between the observations and latent features. However, the downside of model richness is scalability, and notably higher computational costs. Furthermore, such generative models often underperform their competitors if the conditional relationships between model features only partially approximate the underlying generative process (Yu and Deng 2010).

Mixture models are hierarchical models that combine the probability distribution of observation subsets to generate a hybrid distribution. Bayesian mixture models then perform inference by incorporating the mixture likelihood with prior distributions for both the component hyper-parameters and the mixture probabilities. A common extension to Bayesian mixture models is to assume the latent parameters defining the mixture component identities are connected in a Markov chain, rather than independent identically distributed random variables. The resulting models are sequential hierarchical models such as Hidden Markov models, Kalman filters and particle filters (Arulampalam et al. 2002). Another extension are Bayesian nonparametrics (Müller and Quintana 2004), which introduce priors are placed upon a number of components through methods like Dirichlet point processes in order to improve accuracy.

Bayesian Markov-based model for temporal spectrum prediction is essentially a doubly stochastic tracking problem in the time dimension. Within a doubly stochastic model the distribution of the random variable of interest is assumed, or observed, to be a standard, pre-defined distribution and the related parameters of the distribution itself are assumed to be time varying random variables. Through this framework, numerous models have been proposed for modelling aspects of spectrum utilisation, through the use of Poisson distributed Markov chain models proposed by (Bayhan and Alagöz 2012); Hidden Markov Models of (Eltom et al. 2018b); two-state (Csurgai-Horváth and Bito 2011) and three-state discrete-time Markov Models (Bayhan and Alagöz 2012); and higher-order Markov chains (Li et al. 2010).

In contrast to the aforementioned approaches, Gaussian models utilise a Kalman filter, either directly for Linear models, or through a sub-optimal variant for Nonlinear models. These variants include, but are not limited to, the extended Kalman filter, and uncentered Kalman filters. More general nonlinear models can also be constructed using particle filter methods and Monte-Carlo Markov Chain (MCMC) approximations to construct the conditional posterior probability for the relationship between the latent variables and the observations (Kobayashi et al. 2011).

In contrast to the aforementioned stochastic methods, another popular framework for deterministic modelling is the Autoregressive Moving Average (ARMA) model, and its associated techniques. These approaches assume that future behaviour of a time series is a product of a weighted sum of past observations, and an additional set of random contributions that cannot be observed. ARMA is particularly appealing for time series modelling, due to both its simple implementation and easy of interpretation—in contrast to machine

learning models—as the output state is not a black box. ARMA variants have been widely used across a range of time series prediction tasks, ranging from predicting short-term stock price movements (McLeod and Zhang 2008), retail sales behaviour (Ramos et al. 2015), and residential electricity consumption forecasting (Badrinath Krishna et al. 2016). This level of support for the technique is driven by the simplicity of ARMA, on both a conceptual and computational level. This simplicity has allowed ARMA to be generalised to multivariate time series, of the form likely to be seen within the spectrum allocation context, through Vector Autoregression (VAR) and the Vector Autoregression Moving-Average (VARMA) (Lütkepohl 2006).

However, it must be noted that ARMA-like models are limited by their reliance upon regularly sampled data which must also be stationary. The stationary condition imposes that all properties of the modelled process must be constant with time, which is not necessarily true for spectral occupancy, which can have seasonalities and trends on time scales ranging from days to years. To account for this, it has been shown that repeatedly differentiating the time series until the point where the sequence can be considered to be stationary works, producing what is known as an Autoregressive Integrated Moving Average (ARIMA) model. The influence of seasonalities can also be explicitly incorporated through a seasonal ARIMA (SARIMA) (Chatfield 1980, Chap. 4), the idea of which has been extended further by Facebook's Prophet algorithm (Taylor and Letham 2017), which is a variant of Generalised Additive Model (GAM) that attempts to decompose the time series into components corresponding to the trend, the seasonality and the influence of holidays (Harvey and Peters 1990). In contrast to ARIMA-like models, Prophet has been designed to manage irregularly spaced data, subject to the assumption that this data can be assumed to have an underlying continuity. By adding additional features, more non-stationary traffic processes, such as trends in vehicular movement (Williams 2001; Jiang 2002) and freeway traffic speeds (Chandra and Al-Deek 2009) have been able to be modelled using extensions of ARIMA. However, even in these works, the autoregressive core of ARIMA means that the approach is still limits these techniques to data sets in which the mean and variance of the underlying signal do not change with time.

Another concern with ARIMA-like methods is their ability to resolve discontinuities and shocks (Christiano et al. 2006; Chari et al. 2008; Kascha 2012). In a signals context such behaviours are visible at very short time scales as devices cease or begin activity; and in very long scales, where aggregate features change, such as would be the case for observations at a time scale that would not, for example, capture the transition from a city during day- and night-times.

Decision trees are another viable technique for exploring the relationship between observations, and their dependent outputs, in a manner that is suitable for understanding the factors that drive changes in the observed spectral environment. Such trees are hierarchically structured topologies of branches connecting decision points, connected by nodes representing bifurcations in the state after the decision point (Rokach and Maimon 2008). In aggregate, the unfolding set of decision nodes produces a model that can readily be interpreted, in contrast to many other nonlinear techniques.

A powerful approach for constructing regression trees can be found in Gradient boosting, with a variant known as eXtreme Gradient Boosting (XGBoost) (Chen and Guestrin 2016), which produces an ensemble of models that, in aggregate, exhibit strong convergence properties while preserving scalability and minimising over-fitting. The potential for XGBoost to predict spectral occupancy at a receiver is underscored by its success in other, similar domains, including electricity demand (Ben Taieb and Hyndman 2014) and oil price forecasting (Zhou et al. 2019).

## 4.1.2 Machine learning methods for temporal prediction

Artificial Neural Networks (ANNs) are a family of computational techniques for machine learning, that draws upon ideas of mathematical biology. These tools can be used to predict, classify, or cluster data, through what are essentially large systems of nonlinear equations, subject to sets of free parameters. By exposing the model to data, the underlying parameter space can be refined, so that the nonlinear system of equations begins to reflect the data in a manner that allows for the future state of the system to be predicted.

One primary advantage of machine learning type approaches are their ability to incorporate contextual information. This information can include information including, but not limited to, device mobility data, or information from nearby sensors or correlated frequency bands. In doing so, we can generalise Eq. 2 to take the more flexible form

$$O(t, \mathbf{x}_i, b_j, \hat{P}) = f_{i,j}(t; g(O, \mathbf{x}(t), b_j)), \quad (3)$$

where  $g$  represents the fusion of spectral–spatio–temporal contextual data. While the synthesis of any such information into the predictive framework does not guarantee increased prediction accuracy, through careful design it should be possible to improve the generalisability of the constructed predictions.

**4.1.2.1 RNNs** Recurrent Neural Networks (RNNs) are a family of neural networks which are ideally suited to problems involving sequentially structured data. While traditional feed-forward neural networks can capture temporal dynamics (Chen et al. 2021), they are often limited to sequences where that the input to each node must be of a fixed length. In contrast, RNNs encode cyclical structural dependencies, that are more appropriate for capturing the sequential nature of time series data. These cyclic structures also allows for parameter sharing to occur, as a number of distinct nodes will all be subject to the same parameters, allowing for the model to efficiently learn correlations between elements of a sequence. Furthermore, because a node with a self-loop may be unfolded an arbitrary number of times, there is no fundamental restriction placed upon the length of the input sequence to the node. This property is of particular relevance to time series problems, as it allows a single model to be constructed for time series with different lengths.

Beyond these cyclical structural elements, RNNs can also be constructed in fashion that allows them to capture: sequential inputs with non-sequential outputs (or vice versa); cases where the input and output sequence are of different lengths (as seen in encoder-decoder architectures); where there are output-to-hidden loops instead of hidden-to-hidden loops (teacher forcing); and bidirectional structures, in which connections extend both forward and backwards in time.

Due to the flexibility in which RNNs can be constructed, they have been successfully applied to problems with crossover to the task of spectral prediction by (Yasdi 1999; Van Lint et al. 2002), and (Song et al. 2016), who all considered problems stemming from traffic dynamics and human mobility. Other comparable problem domains have included predicting wireless user activity (Agarwal et al. 2016) and detecting anomalous behaviours in radio networks (O’Shea et al. 2016; Katzev et al. 2020, 2021, 2022).

**4.1.2.2 LSTMs** While RNN architectures are able to conduct time series predictions, they often struggle as the time scale of interdependence increases, as the training process struggles to update weights as the sequence length increases (Bengio et al. 1994). To resolve this, Long Short-Term Memory (LSTM) cells have been introduced as hidden units, as an alter-



native to RNN cells. These cells take the structure of an RNN cell, subject to the addition of an additional logical element known as a forget gate (Hochreiter and Schmidhuber 1997). This gate, which is unique to LSTMs, takes the form of a vectors of numbers between 0 and 1 that depend upon the input and previous state through trainable parameters, that focuses the attention on subsets of the previous state. Through the addition of these gates, LSTMs are able to learn which parts of the input and previous state to keep or forget, and to what degree, which allows for modelling complex and arbitrarily long long-term dependencies, beyond what RNNs are capable of.

Because of this, LSTMs have quickly become the most commonly utilised technique for regression and classification tasks in a time series context. Within a spectrum prediction context, several recent works have demonstrated that LSTMs can successfully predict temporal problems from the dynamic spectrum allocation space (Radhakrishnan et al. 2021b), in a fashion that exhibits lower computational complexity than prior approaches (Radhakrishnan and Kandeepan 2020; Radhakrishnan et al. 2021a). When considering the broader set of problem spaces with similar dynamics to those seen within dynamic spectrum allocation, LSTMs have also been shown to be an accurate for predicting nonlinear vehicular traffic dynamics (Zhao et al. 2017b), weather, precipitation (Shi et al. 2015), and energy consumption forecasting (aggregate and single-meter) (Jian et al. 2017a).

LSTMs can also be used in conjunction with what are known as Deep Belief Networks (DBNs), which are a form of Bayesian neural network. Such Bayesian models attempt to predict not just a value, but its probability distribution. These stacked models have been shown to be successful for developing Bayesian nonlinear high dimensional models where other Bayesian models like HMM have struggled (Melchior et al. 2017). For spectrum occupancy modelling, DBN in conjunction with Long Short Term Memory networks have the potential to capture the inherent inter-dependencies between the spatial, temporal and spectral features of the channel states.

Another potential avenue for augmenting LSTMs can be found in Physics-Guided Neural Networks (Karpadne et al. 2017). Such approaches attempt to regularise the behaviour of the solution through a loss function that considers a basic, physical model of the underlying system behaviour. For the purposes of modelling Spectral Occupancy, such a model could be used to increase the effective spatial density of sensor networks, by approximating the distribution of spectral occupancy in their neighbourhoods in terms of known, physics based relationships.

**4.1.2.3 Attention mechanisms** Attention mechanisms (Bahdanau et al. 2014) have recently grown in popularity—especially within the Natural Language Processing community—for their ability to create complex, representative models by calculating the relationship between different positions within sequential data. This mechanism works by constructing a matrix between the input and output sequences, which can be thought of as a learned representation of correlation between each element of the input sequence to each element of the output sequence. Within a Natural Language Processing context, if the input sequence is a sentence in one language, and the output the same sentence translated into another language, then the correlations measure how important an individual word is to the output translation, and how its context also affects the translation.

Through this approach, attention mechanisms are able to better capture long-range dependencies, which are particularly difficult to represent within RNN and LSTM architectures. Similar to RNNs and LSTMs, Attention mechanisms have often been embedded within encoder–decoder architectures (Luong et al. 2015), although recent work has

suggested that Attention mechanisms perform best when embedded in a network architecture known as a Transformer, in which the model behaviour is entirely driven by the Attention mechanism itself (Vaswani et al. 2017).

While the development of Attention mechanisms has primarily been focused upon sequential data, like temporal sequences, recent work has also demonstrated its promise for spatial data (Xu et al. 2015). This in turn suggests that Attention mechanisms may hold promise for the spectral–spatio–temporal aspect of forecasting Spectral Occupancy, either directly or as part of a broader modelling framework. However, it must be noted that the ability for Attention mechanisms to scale with higher–dimensional data remains an open question.

**4.1.2.4 Convolutional networks** One of the most commonly utilised machine learning techniques are Convolutional Neural Networks (CNNs), which utilise convolution operations to leverage structural dependencies in regular data to quickly generate accurate representative models. Unlike the techniques discussed to this point, a CNNs can be thought of as acting not directly upon the data itself, but rather upon a filtered form of said data. In the context of an image, a filter could be thought of as taking a subset of, blurring, or down–sampling the image, or some combination of the aforementioned operations. A neural network that is built in terms of such filtered components reduces the cost of discovering inter–dependency of the individual components of the neural network, while simultaneously improving convergence properties in a manner that minimises the overall computational cost (Goodfellow et al. 2016, Chap. 9) while producing successful results in a range of domains, including predicting wireless interference (Schmidt et al. 2017) and human mobility (Zhang et al. 2016b).

Typically CNNs are employed in image-based problems, due to the inherent regularity and translation invariance of pixel structured data (Zhang et al. 1990). However, in occupancy data similar regularity can be constructed through careful design of sensor networks. Moreover, CNNs have been applied to spatiotemporal processes by simply adding time as an additional matrix dimension, which can introduce underlying structural seasonalities that is well suited to CNNs (Pyrkov et al. 2018), making them potentially appropriate for both sequential and parallel spatiotemporal prediction. Through this, CNNs have been successfully employed in predicting wireless interference (Schmidt et al. 2017) and human mobility (Zhang et al. 2016b).

One promising variant for parallel spectral occupancy prediction can be found in the Wavenet (Oord et al. 2016) variant of CNN, which employs dilated convolutions are applied in an autoregressive manner through the introduction of an offset stride between temporal components. Wavenet uses to date have primarily been in audio prediction, however such data contains structural similarities with constructing predictions across a range of frequency bands and locations across the spectral–spatio–temporal domain.

**4.1.2.5 Graph convolutional networks** In the case where data cannot be expressively discretised upon a regular grid the performance of CNNs begins to degrade. In response to this, a substantial volume of recent research has focused upon extending convolution operators onto less regularly structured data, through the use of graph networks. These networks often utilise spectral graph theory (Bruna et al. 2013), and construct the convolution operator through a Fourier transform on the graph. This Fourier transform is in turn a generalisation of the Laplacian operator to a graph-based Laplacian, and has been used across a broad range of problem domains including medical classification (Parisot et al. 2017), human activity recognition (Shi et al. 2018), and traffic forecasting (Yu et al. 2017; Hermes et al.

2022). Extensions to the concept of Spectral graph networks can be found in localised spectral graph convolution (Bruna et al. 2013), and multi-scale graph convolution (Abu-El-Haija et al. 2018), which allow for computational efficiency improvements by avoiding having to perform large scale Laplacian eigendecomposition.

However, due to the inherently non-local nature of spectral graph theory, interpreting the convolution operations of such networks can be challenging. This is particularly vexing for some problem domains, like traffic or spectral occupancy forecasting, where interpreting the learned spatial dependencies can provide additional insight. An alternate approach to spectral graph theory can be found in spatial convolutional graph networks, which construct localised convolution operations across the nodes of the graph, where the magnitude of the dynamic edge weighting can give insight into the predicted dynamics (Cui et al. 2018). In the context of a deployed multi-sensor environment, such insight could be used to inform the redeployment of sensor nodes, or the prioritisation of communication between nodes.

## 4.2 Off-sensor predictions

Extending from on-sensor to off-sensor predictions greatly complicates the task of finding analogous problem domains. This in turn increases the complexity of assessing the applicability of extant techniques to the problem at hand. The closest problem domain is numerical weather prediction, in which a finite set of weather stations are used to predict future weather across a broader geographic domain. However even in this case the discrete sensors are augmented by satellite and radar imagery, which is able to capture the exact state across significant proportions of the spatial domain of interest.

Other problem domains that require similar spatiotemporal predictions include vehicular traffic forecasting, and both taxi and ambulance demand forecasting. In the case of traffic forecasting, the motivating data-sets are often recordings of motion through traffic loops on urban road networks. In such cases, while there is a geospatial relationship between the traffic loops, the nature of vehicular traffic means that there is no need to generalise the prediction across the whole problem domain. As such, traffic loop data, while sparse in terms of the physical space it represents is a dense measurement of the behaviours of interest, and it is this density that makes it difficult to map solution techniques from this problem domain into a spectral sensing context.

While taxi and ambulance demand forecasting has a clear spatiotemporal context, the sparse nature of demand for these services means that data is aggregated into large geographic regions (ranging in size from several city blocks up to a suburb), with demand being assumed to be homogeneous over these regions of interest. This assumption of homogeneity introduces a problem similar to that observed with numerical weather prediction—that the characteristic distance scales of such a prediction task are potentially going to be far larger than would be needed to resolve spectral behaviour. Furthermore, the sparse and stochastic nature of demand for these services at any location is fundamentally different to the dense and nominally continuous nature of spectral usage. This again limits the direct applicability of techniques derived for taxi and ambulance demand forecasting to the task of spectral occupancy prediction.

As was alluded to previously, constructing these predictions across the geospatial domain can be performed through two distinct modelling approaches—parallel and sequential prediction. In parallel prediction, the spatial and temporal components (and potentially

also the bandwidth) are considered simultaneously, as part of one single aggregate model; while sequential prediction would involve first constructing predictions forward in time, and then using those predictions in a separate, strictly spatial model to generalise the temporal predictions across the spatial domain. While an integrated, parallel prediction framework may seem intuitively appealing, the relative sparsity of appropriate techniques, and the different nature of occupancy as a function of space as compared to time may lead to undesirable outcomes.

In contrast, sequential prediction allows for drawing from a wider suite of techniques and apply them in the context—be it spatial or temporal prediction—that they are best suited to. At the most basic level, this could be simply adding spatial interpolation (in the form of nearest neighbour, linear, or nonlinear predictions) across the discrete sensor locations, in order to generalise the prediction across space. However, the sequential nature of these models introduces modelling complexity, in terms of how validation is performed upon the spatial layer.

To validate a temporal model at discrete locations, a subset of data (in time) is reserved from the training of any learnable parameters, and then the performance of the model is assessed against the reserved model. To assess the accuracy of the consolidated sequential model at predicting spectral accuracy, spatial predictions should be constructed in terms of predicted (not measured) temporal values at a subset  $S$  of the total receivers  $\mathcal{R}$  (such that  $S \subset \mathcal{R}$ ), and then tested against the corresponding measured temporal values at  $\mathcal{R} \setminus S$ . In doing so, the relative accuracy and performance of different sequential spatio-temporal prediction frameworks can be quantitatively assessed and compared.

#### 4.2.1 Data imputation

Treating off-sensor predictions has motivated a number of recent works in data imputation, in which the geospatial nature of data is used to provide contextual information that can be used to predict behaviours away from sensors. Techniques like nearest-neighbour or linear interpolation have successfully produced data reconstructions without introducing any trainable parameters (Batista and Monard 2003; Zhang 2012). While this significantly reduces the computational burden, on a design level this inflexibility limits the level of response that can be performed to predict domain specific features. To improve upon the accuracy of such approaches, trainable ensembles of interpolation has been shown to improve model accuracy across geospatial domains (Oke et al. 2002). This process has been shown to be an accurate framework for extending the geospatial coverage of weather and climate models (Sanderson et al. 2015). Another similar approach can be found in Kriging, due to its ability to better capture both dependencies and uncertainties across spatial data (Yang et al. 2018).

Drawing from statistical tools, data fusion has been successfully applied to predict the likelihood of occupancy (Eltom et al. 2016). Constructed coalition games (Zhou et al. 2017) and Bayesian non-parametric frameworks (Saad et al. 2012) have also been proposed for managing data fusion under a spectrum occupancy framework. The latter of these tracks the signal amplitude of primary users at each sensor node through a nonlinear particle filter, in order to generate a measure of correlated perception.

Monte Carlo methods (Ni et al. 2005) and Probabilistic Principal Component Analysis (Qu et al. 2009) can also be used to reconstruct the spatiotemporal data by assuming that there is some latent distribution that all points are drawn from. However, recently it has been shown that tensor-completion based mechanisms outperform matrix based

mechanisms (Signoretto et al. 2011), leading to several applications of this process to spatiotemporal traffic prediction tasks (Chen et al. 2019). While tensor-completion approaches have been successfully applied to spatiotemporal data imputation, they are best suited to low levels of missing data. As the level of missing data is increased, it has been shown that Loess based Season and Trend Decomposition (Chaudhry et al. 2019) produces the most accurate data reconstruction (Li et al. 2020).

#### 4.2.2 Neural processes

Neural Processes (Garnelo et al. 2018b) are a novel approach that attempts to bridge the space between deterministic neural networks and Bayesian methods (Garnelo et al. 2018a). In the context of functional approximation and regression, while deep neural networks have been shown to be highly accurate, they require access to a large training corpus for effective training, which needs to be repeated for each problem domain. While this can be somewhat mitigated by employing meta-learning, the computational cost is still significant (Erhan et al. 2009). In contrast, Gaussian Processes, while still computationally expensive, can leverage a specified prior distribution to reduce the inference cost. However, this of course relies upon the ability to construct an appropriate and accurate prior, which is a non-trivial task in and of itself (Ghosal et al. 2006). Neural Processes attempt to construct a neural approximation of Gaussian Process for learning a posterior distribution. By doing so, accurate, probabilistic predictions can be constructed even when exposed to only a limited training corpus. One weakness of Neural Processes is their tendency to under fit data, although this has been broadly addressed by subsequent work, which incorporated attention mechanisms into the deterministic and probabilistic predictions (Kim et al. 2019). Inspired by Neural Processes, (De Brouwer et al. 2019) have developed a method that can incorporate sparse, sporadically observed measurements to update predictions. Of particular interest to geospatial modelling of spectral occupancy is that the correlations between sequences are also learned, so an observation in one time series can update the prediction of all other series of interest, in a manner which would be of particular interest in either off-sensor predictions, or predictions for sensors which are not able to record a measurement due to engineering or communications constraints.

## 5 Supporting technologies

The deployment of any successful system for spectral occupancy prediction requires not only the sensors and prediction techniques, but a framework for managing the deployment, processing and management of the data collected by these systems. Another factor of concern is that the successful development of any such system requires data to test and assess prediction and deployment frameworks. This would necessitate an unfeasible delay between the investment in physical deployment, and the actual deployment of the full spectral occupancy system, and may lead to an improper allocation of deployed assets.

To circumvent this, we propose that the development of such systems should be supported by the development of synthetic data sets that accurately replicate real world dynamics. In doing so, all aspects of the measurement and prediction framework can be tested prior to the investment in any physical capital, which should significantly decrease the risks of such a program. As such this chapter will outline the tools and techniques required to support the deployment of predictive systems for spectral occupancy.

## 5.1 Deployment, processing, and data management

As was alluded to in Sect. 3, the ability to manage and process data to inform predictions of spectral occupancy is of crucial importance. Latencies in the delivery of data are of particular concern for both data processing, due to concerns around data consistency and potential overwrites, and the application of predictive techniques. These issues are particularly acute for Machine Learning algorithms, which are typically designed and trained under scenarios where information flows are predictable and consistently synchronous. In an attempt to resolve this concern, Delay Tolerant Machine Learning has recently been proposed, with authors focusing on preventing update issues between when information is provided and gradients are updated (Sun 2016), and the development of schemes that are robust to delays generated by scattered, asynchronous data sources (Mishchenko et al. 2018). However, while such an approach removes concerns relating to data consistency, and inherent latencies, it may also increase the cost of developing and deploying the sensor network, and introduce complexities in communicating domain information outside the domain sampled by an individual sensor.

A complimentary approach may be found in down-sampling data, either by way of low-pass filtering strategies like the Exponential Moving Average (Haque et al. 2015), or employing subsets of more complex decompositions (Wen et al. 2019). In doing so, stored data volumes will decrease, but the down-sampling process itself will likely introduce additional computational costs, and potentially prevent the ability to develop real-time predictions, especially under certain predictive time scales.

Under either an offline or online framework, specialised hardware may be required to efficiently execute Machine Learning algorithms. These computational considerations are particularly important as machine learning typically has a computational cost that is orders of magnitude higher than classical signal processing techniques. As such, any utilisation of Machine Learning must be considered in context of the scalability of computational resources.

The ubiquity of CPU and GPU based systems has made them de rigueur for Machine Learning applications. Of these, GPU based implementations can perform machine learning tasks an order of magnitude faster than the equivalent CPU, at the cost of requiring 8 times as much energy (Zhao et al. 2017a). However, more specialised hardware has begun to gain prominence, with particular focus being placed upon Field-Programmable Gate Arrays (FPGAs) and Application-Specific Integrated Circuits (ASICs) due to their ability to improve computational performance while minimising power consumption.

FPGAs, as their name suggests, can be used in multiple contexts, are re-programmable, and can even have parts of the chip reconfigured while other components are under load. In contrast, ASICs exhibit permanent circuitry, and are designed to be operated within a specific context. The design specificity allows ASICs to be more energy efficient than equivalent FPGA chips, however, this comes at a higher cost due to the inherent difficulties in developing and utilising a chip that has been produced for one specific purpose.

Relative to CPUs, (Zhao et al. 2017a) found that FPGAs could be 2.5 times more powerful than CPUs for CNN inference using binary arithmetic, while simultaneously performing 50 times more calculations per watt of energy. This power can be explained by the ability of FPGAs to exploit pipeline parallelism, which allows for batches of multiple data-points to be considered at once, rather than individually (Lacey et al.

2016). Further advances in computational efficiency have been seen in CNNs (Zhang et al. 2016a), ResNets (Nurvitadhi et al. 2017), LSTMs (Wang et al. 2016), and Random Forrest implementations (Nakahara et al. 2017) when employing an optimised pipeline, which in turn suggests that it may be possible to unlock further performance from FPGAs by considering device specific optimisation. The power of such optimisations is underscored in the work of (Wang et al. 2019), who produced results for LSTM time series predictions using an optimised FPGA that were between 20 and 50% better than previous work by (Ferreira and Fonseca 2016). These advantages may be further enhanced in the case of deep neural network architectures, although performance gains are not guaranteed, as care needs to be taken to manage the power budget and on-chip memory of FPGAs (Han et al. 2017). Recent work has demonstrated that careful optimisation of a single-FPGA design can yield real-time modulation classification with 8  $\mu$ S latency and 488k classification throughput (Tridgell et al. 2020).

Moving from FPGAs to ASIC based implementations can in turn introduce additional computational savings, at the expense of the design flexibility of FPGAs. ASICs have been shown to outperform CPU, GPU, and FPGA implementations in both both deep neural networks (Srinivasan et al. 2019) and Gated Recurrent Units (a variant of RNN) (Nurvitadhi et al. 2016). These performance differences can be attributed to ASICs inherently higher compute density and efficiencies, however it must be noted that newer FPGAs will likely close the efficiency gap. ASICs have also been used to accelerate solutions to dense (Chen et al. 2014) and sparse deep neural networks (Albericio et al. 2016), with such accelerators being able to potentially realise significant speedups, while decreasing energy requirements (Zhang et al. 2016c).

## 5.2 Simulating spectral occupancy

While it is possible to collect the requisite data for predicting Spectral Occupancy, doing so at the scales required to assess the accuracy of predictive regimes would require a significant investment. However, this task is only one part of the overall design process, as once the data is collected tools would still be required to manage the prediction of future states of the spectral environment. Doing so would leave the measurement assets stranded with limited utility, until such time that the remainder of the design process had been completed. An alternative development approach could be to simulate the required data, and use it to inform the development of both the sensor networks, and test prediction regimes.

While such an approach would significantly decrease the initial investment required for such a scheme, it would still present significant hurdles. Any such model must take into account not just the behaviour of devices, but also human mobility. These interactions can quickly become highly complex, as human behaviour can be influenced by the interaction between devices and the spectral environment. As such, this section will cover the major facets of this predictive task, and outline tools that can be used in order to support the development of a comprehensive framework to support Dynamic Spectrum Allocation.

### 5.2.1 Empirical signal propagation models

Urban modelling of signal propagation can be traced back to, at least, the work of (Turin et al. 1972), who conducted experimental and statistical work on multi-path signal propagation of analogue and digital signals. Statistical models of urban environments typically classify the environment in terms of broad metrics that include the topographic variation,



vegetation density, building density and height distribution, and density of regions of open area and water (Ibrahim and Parsons 1983).

Of the extant empirical models, the most commonly employed are Okumura–Hata (Falli 1988), Walfisch–Ikegami (Low 1992a), (Erceg et al. 2001), Longley–Rice (Longley and Rice 1968) models, and Log–Normal (Erceg et al. 1999; Walden and Rowsell 2005; Chrysikos and Kotsopoulos 2013) the features of which are summarised in Table 1, with the latter coming into prominence recently due to its ease of fitting. Of these models, Longley–Rice is unique in that it relies less upon experimental measurements, and more on the incorporation of topographic maps, which are used empirically to estimate the geographic influence on large scale signal attenuation.

While the Okumura–Hata, Walfisch–Ikegami and SUI models are able to model signal propagation, they have limited applicability for considering signals in the modern, urban context, due to the spatial length scales involved and the need to model transmitters located below the height of local buildings, as the observed error is inversely proportional to the ratio between the transmitter height and the building height, so that the error increases as the transmitter height is decreased to the building height (or below the building height) (Low 1992b), as the influence of diffraction around building edges and wall reflections becomes more dominant. These factors are further complicated by the density of transmitters in a modern urban environment.

## 5.2.2 Physics–based signal propagation models

To address these concerns, some simulations have resorted to reduced–order, physics inspired models, including simple radial propagation (where the signal is presumed to exist within a fixed radius  $R$  about the source), free–space and two–ray models. However, while such models are computationally efficient, they fail to account for the influence of local environmental conditions, and as such fail to capture physically realistic channel variability.

To capture the influence of the built environment, Ray–tracing techniques were introduced to model the propagation of individual signals. Such models were initially restricted to the 2D plane (Valenzuela 1993), but have since been extended to 3D environments (Kim et al. 1999). Several authors have implemented ray–tracing models for vehicle–to–vehicle transmission under urban, suburban and rural conditions (Maurer et al. 2004), with (Pilosu et al. 2011) incorporating additional pre–processing in order to reduce the computational cost of calculating the attenuation between connected regions of interest. However, while these approaches improve the fidelity of the simulation, they still are fundamentally limited due to the nonlinear growth in computational complexity with the number of agents, and

**Table 1** Properties of empirical modelling approaches

Model	Range	Freq (GHz)	Transmit	Receive	Motivation	Topography
Okumura–Hata	> 5 km	< 1.5	30–200 m	1–10 m	Urban radio	No & No
Walfisch–Ikegami	0.02–5 km	0.8–2	4–50 m	1–3 m	Urban diffusion	Built
SUI	< 10 km	< 2	15–40 m	2–10 m	Urban WiFi	Built*
Longley–Rice	1–2000 km	0.02–20	n/a	n/a	Topography	Natural
Log–Normal	< 2.2 km	1.9–3.5	10–80 m	10–80 m	Urban shadowing	Built* & Natural*

Abstracted topographic modelling approaches denoted by \*

the number of buildings within the environment. As such, these models become computationally intractable as the domain of interest increases to scales of practical relevance.

To address this, frameworks have been developed for constructing computationally efficient propagation models within dense urban environments, through what is known as Beam-tracing (Sridhara and Bohacek 2007). Such an approach can be considered a computationally efficient variant of Ray-tracing, that still allows for incorporating transmission, diffraction, reflection, scattering, channel-gain, and delay-spread within both interior and exterior environments. Of course, such a model requires knowledge of not just building locations and sizing, but also their interior and exterior materials, and while it is more computationally efficient than Ray-tracing, it still shares the same fundamental computational limitations as the size of the simulation region increases.

When considering large scale simulations, path loss propagation models have typically been constructed using simplified representations of the local geometry. Combining these geometric abstractions with pre-calculated statistical measures can be used to minimise the computational cost of the ensuing calculations. An example of such a formulation can be found in the work of (Cheng et al. 2007), who simulated suburban environments using a Dual-Slope Log Distance formulation for long distance signal propagation, and a Nakagami model (Proakis and Salehi 2001) for small-scale signals. These Nakagami models leverage the Nakagami- $m$  probability distribution (Nakagami 1960) in order to approximate the fading and spatial attenuation of wireless signals (Alouini and Goldsmith 2000). A similar Nakagami implementation was also considered by (Khan et al. 2009) for vehicular ad-hoc networks. This dual scale approach is becoming increasingly popular as a method to reduce the computational complexity of solving signal propagation, and has led to systems that scale linearly with the number of agents (Viriyasitavat et al. 2015). A broad taxonomy of these techniques is outlined in Fig. 5.

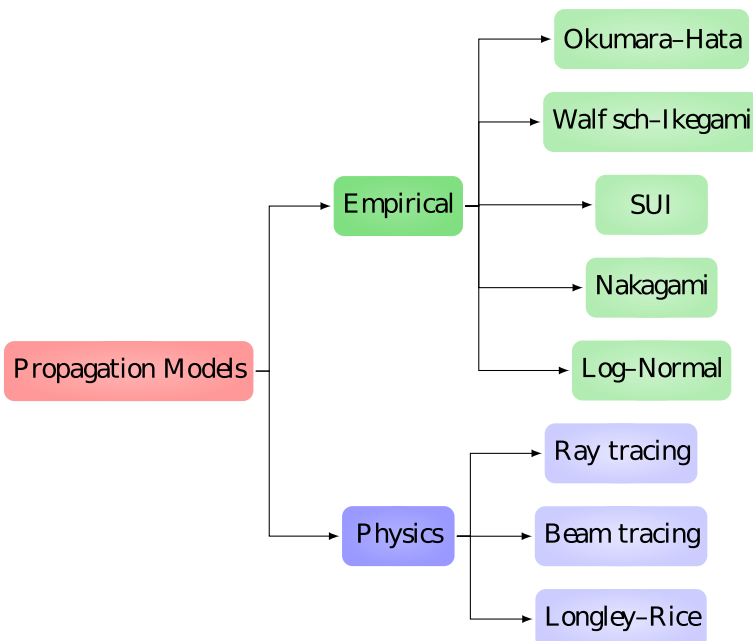


Fig. 5 Classification of spatial propagation models for wireless spectrum

While the geographic scale of the city-scale environments is the primary driver of signal attenuation in urban environments, multiple other factors influence spectral propagation. The nature of built topography is both a secondary driver of attenuation—as well as being a primary source of both shadowing and fading—with other factors including weather effects, especially rain, cloud cover, and smog all influencing signal attenuation through the introduction of additional scattering and diffraction (Jaruwatanadilok et al. 2004). Under ideal circumstances, all of these factors would be accounted for in a model, however there is presently no literature covering a comprehensive signals based model for the built environment that incorporates weather conditions, and as such, the accuracy and applicability of any model must be assessed in the context of the local weather conditions.

### 5.2.3 Device mobility

Beyond the influence of the built and natural environment, a realistic simulacrum of the spectral environment must also take into account device mobility, which in turn is a product of human behaviour and vehicular patterns. While some abstracted, low-resolution data sets do exist (Barlacchi et al. 2015), there is a broad absence of accurate real world data sets that accurately represent large scale spectral-spatio-temporal behaviour at high resolution. This in turn limits the ability for testing techniques for measuring and predicting spectral occupancy requires models to be constructed. Moreover, constructing a representative simulacra of such data is also a non-trivial task, due to the inherently complex and stochastic nature of human behaviour.

The most viable conceptual framework for attempting to construct such a model can be found within Agent Based Modelling (ABM), which attempts to construct rich, representative dynamics not by prescriptive modelling, but rather by establishing motivations for a suite of simulated agents, and then establishing the rules for how these agents can behave and interact. By then iteratively allowing the collective set of agents to evolve to optimise their individual behaviour, rich models full of emergent behaviour can transpire, that have been shown to reflect real world dynamics.

Computationally, the design of an ABM is a competitive process between the number of resolved agents and the spatial and temporal time scales considered, against the extant computational resources. Simulations over limited spatial domains, like public events (Batty et al. 2003) and indoor movement (Crooks et al. 2015) have been resolved to second-level time scales, whereas city-scale simulations are typically modelled over fundamental time scales at the minute (Crooks and Hailegiorgis 2014), hour (Groff 2007), day (Heppenstall et al. 2006), or even year long periods (Haase et al. 2010). By restricting the underlying scales to the smallest individual level of interest, additional resources are freed up for the generation of complexity in the agent dynamics.

Such models can be informed by integrating geographic and population level data, in order to refine the development of the models, and the behaviour of the individual agents. A common source of such data is census information, which can be accessed at resolutions that range from population wide to sub-suburb resolutions. This can also be combined with activity and time-use surveys, such as the UK Time Use Survey (Sullivan and Gershuny 2018) and the American Time Use Survey (ATUS) (Spencer and Aultman-Hall 2019) to inform features like the distribution of work start and end times, the likelihood of people going shopping on work days, and the amount of time spent on key tasks every day. Incorporating such data ensures that the fundamentals of the model are tied in with real world behaviours, without being overly prescriptive.

Further, more granular data can be found in physical footfall surveys, or geolocated and time stamped data from social media networks. Such Geographic Information System (GIS) data sets are particularly interesting, due to their ability to be collected in near-continuous time (Evans 2012). GAMA (GIS & Agent Based Modelling Architecture) and MOSAIC are computational frameworks for merging GIS data with an ABM approach (Taillandier et al. 2019), that has been used for city-scale modelling. More broadly, a Dynamic Data Assimilation (DDA) framework can be employed to refine agent based models in (Ward et al. 2016); and for calibrating dynamical system models of urban crime (Lloyd et al. 2016) and traffic (Work et al. 2008). The fusion of ABMs and DDA has occurred using both Ensemble Kalman Filters (Ward et al. 2016), for which the ensemble is based upon the stochastic evolution of the underlying ABM and the updated data; or Deep Learning Networks (van der Hoog 2017). These data sets can, and should, be used not just for model development, but also to assess the veracity of the ABM as a proxy for real world dynamics. This can be achieved by conducting an exploratory data analysis on validation data sets, that are distinct from the data used for training.

From a computational perspective, scalability is the primary constraint upon the growth of ABMs, due to the nonlinear growth in interactions, and ensuing computational cost (Bazzan and Klügl 2013). As such, city-scale ABMs tend to be built around a small suite of well understood tools, of which MATSim is one of the most popular. This code-base allows for multi-agent simulation of transportation for the purposes of traffic flow management, urban planning, and developing evacuation scenarios, that has been used to simulate more than 50,000 distinct agents at a time (Horni et al. 2016). Multiple extensions of MATSim for city-scale modelling exist, with Tangramob (Castagnari et al. 2018) being one of the more interesting due to its integration of Smart Mobility into the transit model.

For resolving scenarios that rely upon realistic traffic dynamics and interactions, SUMO (Krajzewicz et al. 2002) has proven particularly popular, both for direct simulations, and to manage small-scale traffic interactions within a larger ABM. While traditionally constructed as a 2-dimensional model (Lopez et al. 2018), recent work has shown its viability for contending with 3-dimensional scenarios (Codeca and Härrä 2018). SUMO can be extended further through interaction with JADE, a Java development framework for constructing smart networks (Azevedo et al. 2016). A generalised summary of ABMs and their associated tools can be found in Table 2.

#### 5.2.4 Device use patterns

Of course, for considering spectral occupancy human dynamics models need to be augmented by models for the propagation and use of spectrum contributing devices. Agents within an ABM may contribute to the spectral environment through both static devices, such as personal routers or desktop computers, and dynamic devices like mobile phones and laptops, with multiple devices potentially being attributable to individual agents. These devices will also inherently contribute differently to the spectral environment, based upon differential use cases and power levels.

Simulating the manner through which devices access the spectral environment varies, depending upon the device, and the context of its use. For some parts of the spectrum, like those used by television transmission towers, such access is a constant transmission covering a singular component of the frequency band. For individual device level modelling, probabilistic models have frequently been used to express access demand. Of these, the most commonly utilised approach can be found in memoryless

**Table 2** Comparison of simulation approaches

Name	Language	Usage scenario	Open source	Citation
MATSim	Java	Multi-agent modelling	Yes	Horni et al. (2016)
GAMA	Java	GIS integration	Yes	Taillandier et al. (2019)
SynCity	Java	Energy policy modelling	No	Keirstead et al. (2010)
Polaris	Not available	Traffic	No	Auld et al. (2015)
DEUS	Java	Discrete events	Yes	Picone et al. (2012)
InterSCSimulator	Erlang	Smart cities transit	Yes	Santana et al. (2016)
Sydney Model	Not available	Disruption modelling	No	Cai et al. (2013)
CityFlow-Recon	Not available	Model for GIS integration	No	Li et al. (2017)
SUMO	C++	Traffic Simulation	Yes	Krajzewicz et al. (2002)
LuST	SUMO	Multi-modal traffic simulation	Yes	Codeca et al. (2017)
MoST	SUMO	Multi-modal traffic simulation	Yes	Codeca and Härrri (2018)
Veins	Java, C++	Vehicular Ad-Hoc Networks	Yes	Darus and Bakar (2013)
CityFlow	Python & C++	Traffic Simulation	Yes	Zhang et al. (2019)

Poisson models (Huang et al. 2012), as well as Pareto distribution processes (Neame et al. 1999), Bernoulli processes (Haghani et al. 2007), and Weibull distributions (Lee et al. 2014). These models typically focus upon the distribution of the received time of transmissions or data packets, although they can also be used in the context of aggregated methods, or for specific use contexts.

An alternate approach can be found in Hidden Markov models, which attempt to construct rich simulations of spectral behaviour by considering probabilistic transitions between a finite number of different activity states. A common approach has been to model the transition between on and off states for individual devices (Lee and Seeling 2014). Markov Modulated Fluid models also hold potential, as they have been shown to produce rich dynamics with limited computational cost (Ye et al. 2014). This is made possible by considering events as part of a continuous process in terms of finite flow rates.

The signals transmitted by these devices do not exist in isolation, but rather can be influenced and attenuated by the built and natural environments, and by weather conditions. Such attenuation can be spatially varying, temporally varying, or both. Finally, with the growth of internet enabled devices, the spectral environment can potentially influence the decisions and actions of agents as they move through the environment—for example, by disrupting GPS signals for routing purposes, or interfering with traffic signalling systems. This in turn can create potential behavioural positive or negative feedback-loops, which in the contexts of an ABM could manifest as different forms of emergent behaviour. Another consideration is that the simulation of device level spectral access behaviour must also be flexible enough to account for behavioural changes that would occur under Dynamic Spectrum Allocation, the ramifications of which are currently a topic of active research (Jiang et al. 2016).

## 6 Challenges and future work

While this work has attempted to identify viable research and development pathways for supporting the task of dynamic spectrum allocation, a number of key challenges still remain. To support future work, we have identified some key challenges, that would be well suited for further consideration.

### 6.1 Data features and management

From a data perspective, broad questions still remain about the qualitative and quantitative characteristics of true spectral data at different spatial, temporal and spectral resolutions. Understanding these features has fundamental implications for simulating, measuring, and predicting spectral data. Extracting these features is of particular importance for choosing appropriate predictive frameworks, as it is unlikely that any one single predictive framework would be optimal for all potential signal characteristics. Even the ability to identify such signal characteristics may prove tricky, in the face of potential legal ramifications of conducting large scale signals measurement, storage, and post-processing. However, if that level of data collection is achievable, then one promising avenue for extracting such features algorithmically is motif identification (Alaee et al. 2020), which has been successfully applied to characterise features of time series data. However, additional work would be required to extend such works to the spectral and spatiotemporal domains.

The ability for qualitative and quantitative data features to be characterised could then in turn be used as the basis of either a reference data-set or data set generator that exhibits the characteristics of true signals data that reflects that built and natural environment. The construction of such data would allow for reliable, repeatedly experiments to be performed across multiple works.

It is highly likely that any predictive system, especially one in which sensors communicate with each other would face bottlenecks stemming from the size of the underlying data. As such, it is highly likely that any predictions would be built upon down-sampled or compressed data measurements. This process could be achieved through filtering mechanisms, smoothing, or sub-sampling mechanisms, or in an adaptive fashion by incorporating encoder-style neural networks. Of these, filtering mechanisms have a long established record for facilitating extracting relevant features in temporal data (Baxter and King 1999; Wen et al. 2019).

A further complication on this front is the potential for the recorded data to contain personally identifying features, which may introduce legal complications. This would necessitate a process of data anonymisation, which could draw upon Differential Privacy (Dwork et al. 2006) to ensure that private and identifying information about individual users could not be revealed. However additional research and care would be required to ensure that the temporally continuous nature of the underlying data does not give rise to opportunities for re-identification of individual signal sources (Culnane et al. 2019).

Any consideration of data management, aggregation, anonymisation, and down-sampling must, of course, be assessed in the context of how these changes impact upon the predictability of the underlying signals. While many of these approaches should reduce the inherent variability of the underlying signals, making them easier to predict, these changes may well be offset by the added data variance of any anonymisation process. Moreover, there is the potential for these changes to introduce systemic biases in the ability to predict

and understand the state of the spectral environment, which also needs to be considered in any subsequent research.

## 6.2 Adversarial interactions

One important consideration when considering deployed systems is the potential for adversarial interactions. This is especially true given the scale and socioeconomic value of the spectral environment, which would introduce incentives for malicious agents to attempt to attack any deployed systems. Such adversarial behaviours have been repeatedly studied within an image-based context (Biggio et al. 2013), with recent works also beginning to consider adversarial sensitivity in a time series context (Karim et al. 2020). As such open questions still remain about both the potential to exploit large scale time-series models, and the availability of mechanisms to deter such adversarial behaviour. Of course, the presence of adversarial attacks also opens the potential for supporting dynamic spectrum allocation with appropriate adversarial defences, or certifications against adversarial behaviour (Lecuyer et al. 2019; Cullen et al. 2022, 2023).

## 6.3 Anomaly detection and fault tolerance

While anomaly detection has been well studied—including relevant works relating to time-series (Laptev et al. 2015)—there is still scope for further work in developing systems that are reactive and responsive to anomalous behaviours. In a dynamic spectrum allocation context, an anomaly could present from a component in the sensor network failing; or it could be the product of a large scale grouping of people, like at a sporting event. As such there is a need for studying how tools and processes relating to both anomaly detection and fault tolerant prediction can be applied to generate predictions that are accurate in the face of both genuine faults and outlier events.

## 7 Conclusion

Growth of demand for wireless data, and the commensurate limitations on the amount of spectral resources that can be allocated, are quickly growing to be significant problems worldwide. Servicing that demand necessitates a change in the manner through which devices access and understand the spectral environment. The high dimensionality of the problem space, and the difficulty in accessing the true state makes this problem particularly well suited for exploring through machine learning techniques.

To support the engineering challenge of building deployable systems, a range of prediction techniques has been reviewed, and assessed for their applicability to different facets of the prediction task. This has been performed with the aim of constructing accurate predictions of spectral occupancy across large, heterogeneous geographic environments. This included a study of state of the art machine learning techniques that could be deployed to facilitate this task, and how a distributed sensing framework could be deployed to achieve this. Also included in this review was a consideration of how realistic data could be constructed, as a proxy to measuring the true spectral environment, in order to test how various techniques performed prior to implementation.

While the presented techniques have the potential to inform spectral management and prediction, the deployment of such a system would also need to consider the economics of



a dynamic spectrum environment. Such work would need to consider not just how to price spectrum access, but broader questions of the relative merits of cooperative environments or competitive market mechanisms, and how such environments would affect the motivations of adversarial actors. We believe that resolving such questions is the next frontier for developing and implementing a cooperative, Dynamic Spectrum Allocation environment that will scale with future spectral demand.

**Acknowledgements** The authors would like to extend their thanks to David Boland of the University of Sydney; Marc Katzef and Gregory J. Karanikas of the University of Melbourne; Daniel Salmond and Perry Plackmore from the Defence Science and Technology Group for their assistance and knowledge in exploring this problem domain. The research within this paper received funding from the Australian Government through Trusted Autonomous Systems, a Defence Cooperative Research Centre funded through the Next Generation Technologies Fund.

**Funding** Open Access funding enabled and organized by CAUL and its Member Institutions. The research within this paper received funding from the Australian Government through Trusted Autonomous Systems, a Defence Cooperative Research Centre funded through the Next Generation Technologies Fund.

**Data availability** Not applicable.

**Code availability** Not applicable.

## Declarations

**Conflict of interest** Beyond the previously declared funding, the authors have no applicable conflicts of interest to declare.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Abo-Zahhad M, Sabor N, Sasaki S et al (2016) A centralized immune-Voronoi deployment algorithm for coverage maximization and energy conservation in mobile wireless sensor networks. *Inf Fusion* 30:36–51
- Abu-El-Hajja S, Kapoor A, Perozzi B, et al (2018) N-GCN: multi-scale graph convolution for semi-supervised node classification. arXiv preprint [arXiv:1802.08888](https://arxiv.org/abs/1802.08888)
- Agarwal A, Dubey S, Khan MA, et al (2016) Learning based primary user activity prediction in cognitive radio networks for efficient dynamic spectrum access. In: 2016 international conference on signal processing and communications (SPCOM), IEEE, pp 1–5
- Ahmed N, Kanhere SS, Jha S (2005) The holes problem in wireless sensor networks: a survey. *ACM SIGMOBILE Mob Comput Commun Rev* 9(2):4–18
- Akkaya K, Younis M (2005) A survey on routing protocols for wireless sensor networks. *Ad Hoc Netw* 3(3):325–349
- Al-Fuqaha A, Guizani M, Mohammadi M et al (2015) Internet of things: a survey on enabling technologies, protocols, and applications. *IEEE Commun Surv Tutor* 17(4):2347–2376
- Alaee S, Kamgar K, Keogh E (2020) Matrix profile XXII: exact discovery of time series motifs under DTW. In: 2020 IEEE international conference on data mining (ICDM), IEEE, pp 900–905
- Albericio J, Judd P, Hetherington T et al (2016) CNVLUTIN: ineffectual-neuron-free deep neural network computing. *ACM SIGARCH Comput Archit News* 44(3):1–13

- Ali A, Hamouda W (2016) Advances on spectrum sensing for cognitive radio networks: theory and applications. *IEEE Commun Surv Tutor* 19(2):1277–1304
- Alouini MS, Goldsmith AJ (2000) Adaptive modulation over Nakagami fading channels. *Wirel Pers Commun* 13(1–2):119–143
- Arulampalam MS, Maskell S, Gordon N et al (2002) A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Trans Signal Process* 50(2):174–188
- Auld J, Hope M, Ley H et al (2015) POLARIS: Agent-based modeling framework development and implementation for integrated travel demand and network and operations simulations. *Transp Res C*. <https://doi.org/10.1016/j.trc.2015.07.017>
- Azevedo T, De Araújo PJ, Rossetti RJ, et al (2016) JADE, TraSMAP and SUMO: a tool-chain for simulating traffic light control. [arXiv:1601.08154](https://arxiv.org/abs/1601.08154)
- Badrinath Krishna V, Iyer RK, Sanders WH (2016) ARIMA-based modeling and validation of consumption readings in power grids. In: Rome E, Theocharidou M, Wolthusen S (eds) *Critical information infrastructures security. Lecture notes in computer science*. Springer, Berlin, pp 199–210
- Bahdanau D, Cho K, Bengio Y (2014) Neural machine translation by jointly learning to align and translate. [arXiv preprint arXiv:1409.0473](https://arxiv.org/abs/1409.0473)
- Barlacchi G, De Nadai M, Larcher R et al (2015) A multi-source dataset of urban life in the city of Milan and the province of Trentino. *Sci Data* 2(1):1–15
- Batista GE, Monard MC (2003) An analysis of four missing data treatment methods for supervised learning. *Appl Artif Intell* 17(5–6):519–533
- Batty M, Desyllas J, Duxbury E (2003) Safety in numbers? Modelling crowds and designing control for the Notting hill carnival. *Urban Stud* 40(8):1573–1590
- Baxter M, King RG (1999) Measuring business cycles: approximate band-pass filters for economic time series. *Rev Econ Stat* 81(4):575–593
- Bayhan S, Alagöz F (2012) Distributed channel selection in CRAHNs: a non-selfish scheme for mitigating spectrum fragmentation. *Ad Hoc Netw* 10(5):774–788
- Bazzan A, Klügl F (2013) A review on agent-based technology for traffic and transportation. *Knowl Eng Rev* 29:375–403. <https://doi.org/10.1017/S0269888913000118>
- Ben Taieb S, Hyndman RJ (2014) A gradient boosting approach to the Kaggle load forecasting competition. *Int J Forecast* 30(2):382–394. <https://doi.org/10.1016/j.ijforecast.2013.07.005>
- Bengio Y, Simard P, Frasconi P (1994) Learning long-term dependencies with gradient descent is difficult. *IEEE Trans Neural Netw* 5(2):157–166. <https://doi.org/10.1109/72.279181>
- Biggio B, Corona I, Maiorca D et al (2013) Evasion attacks against machine learning at test time. In: *Joint European conference on machine learning and knowledge discovery in databases*. Springer, Berlin, pp 387–402
- Blum RS, Kassam SA, Poor HV (1997) Distributed detection with multiple sensors II. Advanced topics. *Proc IEEE* 85(1):64–79
- Bruna J, Zaremba W, Szlam A, et al (2013) Spectral networks and locally connected networks on graphs. [arXiv preprint arXiv:1312.6203](https://arxiv.org/abs/1312.6203)
- Cai C, Wang Y, Geers G (2013) Time-variant travel time prediction model and its application in a cooperative traffic control system. *Transp Res Rec* 2381(1):36–44
- Candes EJ, Romberg JK, Tao T (2006) Stable signal recovery from incomplete and inaccurate measurements. *Commun Pure Appl Math* 59(8):1207–1223
- Castagnari C, Corradini F, De Angelis F, et al (2018) Tangramob: an agent-based simulation framework for validating urban smart mobility solutions. [arXiv:1805.10906](https://arxiv.org/abs/1805.10906)
- Chandra SR, Al-Deek H (2009) Predictions of freeway traffic speeds and volumes using vector autoregressive models. *J Intell Transp Syst* 13(2):53–72
- Chang CY, Chang CT, Chen YC et al (2008) Obstacle-resistant deployment algorithms for wireless sensor networks. *IEEE Trans Veh Technol* 58(6):2925–2941
- Chari VV, Kehoe PJ, McGrattan ER (2008) Are structural VARs with long-run restrictions useful in developing business cycle theory? *J Monet Econ* 55(8):1337–1352
- Chatfield C (1980) *The analysis of time series: an introduction*. Chapman and Hall; Chapman and Hall in association with Methuen, London; New York, N.Y., oCLC: 682061640
- Chaudhry A, Li W, Basri A et al (2019) A Method for Improving Imputation and Prediction Accuracy of Hhighly Seasonal Univariate Data with Large Periods of Missingness. *Wirel Commun Mob Comput*. <https://doi.org/10.1155/2019/4039758>
- Chen T, Guestrin C (2016) XGBoost: a scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. ACM, pp 785–794. <https://doi.org/10.1145/2939672.2939785>

- Chen Y, Luo T, Liu S, et al (2014) DaDianNao: a machine-learning supercomputer. In: Proceedings of the 47th annual IEEE/ACM international symposium on microarchitecture, IEEE Computer Society, pp 609–622
- Chen X, He Z, Sun L (2019) A Bayesian tensor decomposition approach for spatiotemporal traffic data imputation. *Transp Res C* 98:73–84
- Chen J, Yang S, Zhang D et al (2021) A turning point prediction method of stock price based on RVFL-GMDH and Chaotic Time Series Analysis. *Knowl Inf Syst* 63(10):2693–2718
- Cheng L, Henty BE, Stancil DD et al (2007) Mobile vehicle-to-vehicle narrow-band channel measurement and characterization of the 5.9 GHz dedicated short range communication (DSRC) frequency band. *IEEE J Sel Areas Commun* 25(8):1501–1516
- Christiano LJ, Eichenbaum M, Vigfusson R et al (2006) Assessing structural VARs. *NBER Macroecon Annu* 21:1–105
- Chrysikos T, Kotsopoulos S (2013) Site-specific validation of path loss models and large-scale fading characterization for a complex urban propagation topology at 2.4 GHz. In: Proceedings of the international multiconference of engineers and computer scientists, pp 2078–0958
- Codeca L, Härrä J (2018) Monaco SUMO traffic (MoST) scenario: a 3D mobility scenario for cooperative ITS. In: SUMO 2018, SUMO user conference, simulating autonomous and intermodal transport systems
- Codeca L, Frank R, Faye S et al (2017) Luxembourg SUMO traffic (LuST) scenario: traffic demand evaluation. *IEEE Intell Transp Syst Mag* 9(2):52–63
- Cordeiro C, Challapali K, Birru D, et al (2005) IEEE 802.22: the first worldwide wireless standard based on cognitive radios. In: First IEEE international symposium on new frontiers in dynamic spectrum access networks, 2005. DySPAN 2005, IEEE, pp 328–337
- Corke P, Hrabar S, Peterson R, et al (2004) Autonomous deployment and repair of a sensor network using an unmanned aerial vehicle. In: IEEE international conference on robotics and automation, 2004. Proceedings. ICRA'04. 2004, IEEE, pp 3602–3608
- Crooks AT, Hailegiorgis AB (2014) An agent-based modeling approach applied to the spread of cholera. *Environ Model Softw* 62:164–177
- Crooks A, Croitoru A, Lu X et al (2015) Walk this way: improving pedestrian agent-based models through scene activity analysis. *ISPRS Int J Geo-Inf* 4(3):1627–1656
- Csurgai-Horváth L, Bito J (2011) Primary and secondary user activity models for cognitive wireless network. In: Proceedings of the 2011 11th international conference on telecommunications (ConTEL), pp 301–306
- Cui Z, Henrickson K, Ke R, et al (2018) Traffic graph convolutional recurrent neural network: a deep learning framework for network-scale traffic learning and forecasting. arXiv preprint [arXiv:1802.07007](https://arxiv.org/abs/1802.07007)
- Cullen AC, Montague P, Liu S, et al (2022) Double bubble, toil and trouble: enhancing certified robustness through transitivity. arXiv preprint [arXiv:2210.06077](https://arxiv.org/abs/2210.06077)
- Cullen AC, Montague P, Liu S, et al (2023) Exploiting certified defences to attack randomised smoothing. [arXiv:2302.04379](https://arxiv.org/abs/2302.04379) [cs]
- Culnane D, Rubinstein A, Benjamin I, et al (2019) Stop the open data bus, we want to get off. arXiv preprint [arXiv:1908.05004](https://arxiv.org/abs/1908.05004)
- Darus MY, Bakar KA (2013) Congestion control algorithm in VANETs. *World Appl Sci J* 21(7):1057–1061
- De Brouwer E, Simm J, Arany A, et al (2019) GRU-ODE-Bayes: continuous modeling of sporadically-observed time series. In: Advances in neural information processing systems, pp 7379–7390
- Do J, Akos DM, Enge PK (2004) L and S bands spectrum survey in the San Francisco Bay Area. In: PLANS 2004. Position location and navigation symposium (IEEE Cat. No. 04CH37556), IEEE, pp 566–572
- Dwork C, McSherry F, Nissim K et al (2006) Calibrating noise to sensitivity in private data analysis. In: Theory of cryptography conference. Springer, Berlin, pp 265–284
- El-Khamy SE, El-Mahallawy MS, Youssef ENS (2013) Improved wideband spectrum sensing techniques using wavelet-based edge detection for cognitive radio. In: 2013 international conference on computing, networking and communications (ICNC), IEEE, pp 418–423
- Eltom H, Kandeepan S, Liang YC, et al (2016) HMM based cooperative spectrum occupancy prediction using hard fusion. In: 2016 IEEE international conference on communications workshops (ICC), IEEE, pp 669–675
- Eltom H, Kandeepan S, Evans RJ et al (2018a) Statistical spectrum occupancy prediction for dynamic spectrum access: a classification. *EURASIP J Wirel Commun Netw* 1:29
- Eltom H, Kandeepan S, Liang YC et al (2018b) Cooperative soft fusion for HMM-based spectrum occupancy prediction. *IEEE Commun Lett* 22(10):2144–2147
- Erceg V, Greenstein LJ, Tjandra SY et al (1999) An empirically based path loss model for wireless channels in suburban environments. *IEEE J Sel Areas Commun* 17(7):1205–1211

- Erceg V et al (2001) Channel models for fixed wireless applications. IEEE 802.16.3c-01/29r4. Available as [www.ieee802.org/16/tg3/contrib/802163c-01\\_29r4.pdf](http://www.ieee802.org/16/tg3/contrib/802163c-01_29r4.pdf)
- Erhan D, Manzagol PA, Bengio Y, et al (2009) The difficulty of training deep architectures and the effect of unsupervised pre-training. In: Artificial intelligence and statistics, pp 153–160
- Evans A (2012) Uncertainty and error. In: Agent-based models of geographical systems. Springer, pp 37–45
- Falli M (1988) COST 207: digital land mobile radio communications. Tech. Rep. COST 207, Vienna
- Ferreira JC, Fonseca J (2016) An FPGA implementation of a long short-term memory neural network. In: 2016 International Conference on ReConFigurable Computing and FPGAs (ReConFig), IEEE, pp 1–8
- Gardner Jr ES (2006) Exponential smoothing: the state of the art-part II. *Int J Forecast* 22(4):637–666
- Garnelo M, Rosenbaum D, Maddison CJ, et al (2018a) Conditional neural processes. arXiv preprint [arXiv:1807.01613](https://arxiv.org/abs/1807.01613)
- Garnelo M, Schwarz J, Rosenbaum D, et al (2018b) Neural processes. arXiv preprint [arXiv:1807.01622](https://arxiv.org/abs/1807.01622)
- Ghosal S, Roy A et al (2006) Posterior consistency of Gaussian process prior for nonparametric binary regression. *Ann Stat* 34(5):2413–2429
- Goodfellow I, Bengio Y, Courville A (2016) Deep learning. MIT Press, Cambridge
- Gorcin A, Qaraqe KA, Celebi H, et al (2010) An adaptive threshold method for spectrum sensing in multi-channel cognitive radio networks. In: 2010 17th international conference on telecommunications, IEEE, pp 425–429
- Groff ER (2007) Simulation for theory testing and experimentation: an example using routine activity theory and street robbery. *J Quant Criminol* 23(2):75–103
- Haase D, Lautenbach S, Seppelt R (2010) Modeling and simulating residential mobility in a shrinking city using an agent-based approach. *Environ Model Softw* 25(10):1225–1240
- Haghani E, De S, Ansari N (2007) On modeling VoIP traffic in broadband networks. In: IEEE GLOBECOM 2007-IEEE global telecommunications conference, IEEE, pp 1922–1926
- Han S, Kang J, Mao H, et al (2017) ESE: efficient speech recognition engine with sparse LSTM on FPGA. In: Proceedings of the 2017 ACM/SIGDA international symposium on field-programmable gate arrays, ACM, pp 75–84
- Haque ME, Khan MNS, Sheikh MRI (2015) Smoothing control of wind farm output fluctuations by proposed low pass filter, and moving averages. In: 2015 international conference on electrical & electronic engineering (ICEEE), IEEE, pp 121–124
- Harvey AE, Peters SD (1990) Estimation procedures for structural time series models. *J Forecast* 9:89–108. <https://doi.org/10.1002/for.3980090203>
- Hashim HA, Ayinde BO, Abido MA (2016) Optimal placement of relay nodes in wireless sensor network using artificial bee colony algorithm. *J Netw Comput Appl* 64:239–248
- Heo N, Varshney PK (2003) A distributed self spreading algorithm for mobile wireless sensor networks. In: 2003 IEEE wireless communications and networking, 2003. WCNC 2003., IEEE, pp 1597–1602
- Heppenstall A, Evans A, Birkin M (2006) Using hybrid agent-based systems to model spatially-influenced retail markets. *J Artif Soc Soc Simul* 9(3):2
- Hermes L, Hammer B, Melnik A, et al (2022) A graph-based U-net model for predicting traffic in unseen cities. arXiv preprint [arXiv:2202.06725](https://arxiv.org/abs/2202.06725)
- Hochreiter S, Schmidhuber J (1997) Long short-term memory. *Neural Comput* 9(8):1735–1780
- Horni A, Nagel K, Axhausen KW (2016) The multi-agent transport simulation MATSim. Ubiquity Press, London
- Huang CF, Tseng YC (2005) The coverage problem in a wireless sensor network. *Mob Netw Appl* 10(4):519–528
- Huang Y, Wang W, Zhang X, et al (2012) An energy efficient multicast transmission scheme with patching stream exploiting user behavior in wireless networks. In: 2012 IEEE global communications conference (GLOBECOM), IEEE, pp 3537–3541
- Ibrahim M, Parsons J (1983) Signal strength prediction in built-up areas. Part 1: median signal strength. In: IEEE proceedings F (Communications, Radar and Signal Processing), IET, pp 377–384
- Iii JM (2000) An integrated agent architecture for software defined radio. PhD thesis
- Jacques L, Laska JN, Boufounos PT et al (2013) Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Trans Inf Theory* 59(4):2082–2102
- Jaruwatanadilok S, Ketprom U, Kuga Y et al (2004) Modeling the point-to-point wireless communication channel under the adverse weather conditions. *IEICE Trans Electron* 87(9):1455–1462
- Jiang Y (2002) Dynamic prediction of traffic flow and congestion at freeway construction zones. *J Constr Eng* 7(1):45–57
- Jiang D, Ying X, Han Y et al (2016) Collaborative multi-hop routing in cognitive wireless networks. *Wirel Pers Commun* 86(2):901–923

- Jian Z, Cencen X, Ziang Z et al (2017a) Electric load forecasting in smart grids using long-short-term-memory based recurrent neural network. In: 2017 51st annual conference on information sciences and systems (CISS), pp 1–6. <https://doi.org/10.1109/CISS.2017.7926112>
- Jiang C, Zhang H, Ren Y et al (2017b) Machine learning paradigms for next-generation wireless networks. *IEEE Wirel Commun* 24(2):98–105. <https://doi.org/10.1109/MWC.2016.1500356WC>
- Kalita HK, Kar A (2009) Wireless sensor network security analysis. *Int J Next-Gener Netw (IJNGN)* 1(1):1–10
- Karim F, Majumdar S, Darabi H (2020) Adversarial attacks on time series. *IEEE Trans Pattern Anal Mach Intell* 43(10):3309–3320
- Karlof C, Wagner D (2003) Secure routing in wireless sensor networks: attacks and countermeasures. In: Proceedings of the first IEEE international workshop on sensor network protocols and applications, 2003., IEEE, pp 113–127
- Karpatne A, Watkins W, Read J, et al (2017) Physics-guided neural networks (PGNN): an application in lake temperature modeling. arXiv preprint [arXiv:1710.11431](https://arxiv.org/abs/1710.11431)
- Kascha C (2012) A comparison of estimation methods for vector autoregressive moving-average models. *Econom Rev* 31(3):297–324
- Katzef M, Cullen AC, Alpcan T, et al (2020) Distributed generative adversarial networks for anomaly detection. In: International conference on decision and game theory for security, Springer, pp 3–22
- Katzef M, Cullen AC, Alpcan T, et al (2021) Privacy-preserving collaborative SDR networks for anomaly detection. In: ICC 2021-IEEE international conference on communications, IEEE, pp 1–6
- Katzef M, Cullen AC, Alpcan T et al (2022) Generative adversarial networks for anomaly detection on decentralised data. *Annu Rev Control* 53:329–337
- Keirstead J, Samsatli N, Shah N et al (2010) SynCity: an integrated tool kit for urban energy systems modelling. In: Assessment tools and benchmarking practices, energy efficient cities, pp 21–42
- Khalesian M, Delavar MR (2016) Wireless sensors deployment optimization using a constrained pareto-based multi-objective evolutionary approach. *Eng Appl Artif Intell* 53:126–139
- Khan A, Sadhu S, Yeleswarapu M (2009) A comparative analysis of DSRC and 802.11 over vehicular ad hoc networks. Project Report, University of California, Santa Barbara, pp 1–8
- Kim M, Takada J (2009) Efficient multi-channel wideband spectrum sensing technique using filter bank. In: 2009 IEEE 20th international symposium on personal, indoor and mobile radio communications, IEEE, pp 1014–1018
- Kim SC, Guarino B, Willis T et al (1999) Radio propagation measurements and prediction using three-dimensional ray tracing in urban environments at 908 MHz and 1.9 GHz. *IEEE Trans Veh Technol* 48(3):931–946
- Kim H, Mnih A, Schwarz J, et al (2019) Attentive neural processes. arXiv preprint [arXiv:1901.05761](https://arxiv.org/abs/1901.05761)
- Kobayashi H, Mark BL, Turin W (2011) Probability, random processes, and statistical analysis: applications to communications, signal processing. Queueing Theory and Mathematical Finance. Cambridge University Press, New York
- Krajzewicz D, Hertkorn G, Rössel C, et al (2002) SUMO (simulation of urban mobility)-an open-source traffic simulation. In: Proceedings of the 4th middle east symposium on simulation and modelling (MESM20002), pp 183–187
- Lacey G, Taylor GW, Areibi S (2016) Deep learning on FPGAs: past, present, and future. arXiv preprint [arXiv:1602.04283](https://arxiv.org/abs/1602.04283)
- Laptev N, Amizadeh S, Flint I (2015) Generic and scalable framework for automated time-series anomaly detection. In: Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining, pp 1939–1947
- Laska JN, Wen Z, Yin W et al (2011) Trust, but verify: fast and accurate signal recovery from 1-bit compressive measurements. *IEEE Trans Signal Process* 59(11):5289–5301
- Lecuyer M, Atlidakis V, Geambasu R, et al (2019) Certified robustness to adversarial examples with differential privacy. In: 2019 IEEE symposium on security and privacy (SP), IEEE, pp 656–672
- Lee J, Seeling P (2014) Mobile device-level data modeling through high utilization mobile applications. In: 2014 IEEE 11th consumer communications and networking conference (CCNC), IEEE, pp 513–514
- Lee D, Zhou S, Zhong X et al (2014) Spatial modeling of the traffic density in cellular networks. *IEEE Wirel Commun* 21(1):80–88
- Li Y, Dong YN, Zhang H, et al (2010) Spectrum usage prediction based on high-order markov model for cognitive radio networks. In: IEEE 10th international conference on computer and information technology (CIT), pp 2784–2788. <https://doi.org/10.1109/cit.2010.464>
- Li W, Wolinski D, Lin MC (2017) City-scale traffic animation using statistical learning and metamodel-based optimization. *ACM Trans Graph (TOG)* 36(6):200

- Li Z, Xu W, Zhang X et al (2018) A survey on one-bit compressed sensing: theory and applications. *Front Comput Sci* 12(2):217–230
- Li H, Li M, Lin X et al (2020) A spatiotemporal approach for traffic data imputation with complicated missing patterns. *Transp Res C* 119(102):730
- Liang Q, Ren Q (2005) Energy and mobility aware geographical multipath routing for wireless sensor networks. In: IEEE wireless communications and networking conference, 2005, IEEE, pp 1867–1871
- Liao WH, Kao Y, Li YS (2011) A sensor deployment approach using glowworm swarm optimization algorithm in wireless sensor networks. *Expert Syst Appl* 38(10):12180–12188
- Liu X, He D (2014) Ant colony optimization with greedy migration mechanism for node deployment in wireless sensor networks. *J Netw Comput Appl* 39:310–318
- Lloyd D, Santitissadeekorn N, Short MB (2016) Exploring data assimilation and forecasting issues for an urban crime model. *Eur J Appl Math* 27(3):451–478
- Longley AG, Rice PL (1968) Prediction of tropospheric radio transmission loss over irregular terrain. A computer method. Tech. rep, Institute For Telecommunications Science, Boulder CO
- Lopez PA, Behrisch M, Bieker-Walz L, et al (2018) Microscopic traffic simulation using SUMO. In: 2018 21st international conference on intelligent transportation systems (ITSC), IEEE, pp 2575–2582
- Low K (1992a) Comparison of CW-measurements performed in darmstadt with the flat edge model. Tech. Rep. COST 231 TD(92) 8, Vienna
- Low K (1992b) Comparison of urban propagation models with CW-measurements. In: [1992 proceedings] vehicular technology society 42nd VTS conference—frontiers of technology, IEEE, pp 936–942
- Lu H, Li J, Guizani M (2013) Secure and efficient data transmission for cluster-based wireless sensor networks. *IEEE Trans Parallel Distrib Syst* 25(3):750–761
- Lu Q, Yang S, Liu F (2017) Wideband spectrum sensing based on Riemannian distance for cognitive radio networks. *Sensors* 17(4):661
- Luong MT, Pham H, Manning CD (2015) Effective approaches to attention-based neural machine translation. arXiv preprint [arXiv:1508.04025](https://arxiv.org/abs/1508.04025)
- Lütkepohl H (2006) Forecasting with VARMA models. *Handb Econ Forecast* 1:287–325
- Ma Y, Gao Y, Liang YC et al (2016) Reliable and efficient sub-Nyquist wideband spectrum sensing in cooperative cognitive radio networks. *IEEE J Sel Areas Commun* 34(10):2750–2762
- Marĳian A, Vlădeanu C, Marcu I et al (2010) Evaluation of spectrum occupancy in an urban environment in a cognitive radio context. *Int J Adv Telecommun* 3(3):172–181
- Maurer J, Fugen T, Schafer T, et al (2004) A new inter-vehicle communications (IVC) channel model. In: IEEE 60th vehicular technology conference, 2004. VTC2004-Fall, 2004, IEEE, pp 9–13
- McLeod AI, Zhang Y (2008) Faster ARMA maximum likelihood estimation. *Comput Stat Data Anal* 52(4):2166–2176. <https://doi.org/10.1016/j.csda.2007.07.020>. [arXiv:1611.00965](https://arxiv.org/abs/1611.00965)
- Melchior J, Wang N, Wiskott L (2017) Gaussian-binary restricted Boltzmann machines for modeling natural image statistics. *PLoS ONE* 12(2):e0171015
- Mishchenko K, Iutzeler F, Mallick J, et al (2018) A delay-tolerant proximal-gradient algorithm for distributed learning. In: International conference on machine learning, pp 3584–3592
- Müller P, Quintana FA (2004) Nonparametric Bayesian data analysis. *Stat Sci* 19:95–110
- Nakagami N (1960) The m-distribution, a general formula of intensity of rapid fading. In: Statistical methods in radio wave propagation: proc. of a symp. held June 18–20, 1958, Pergamon Press, Oxford
- Nakahara H, Jinguji A, Sato S, et al (2017) A random forest using a multi-valued decision diagram on an FPGA. In: 2017 IEEE 47th international symposium on multiple-valued logic (ISMVL), IEEE, pp 266–271
- Neame TD, Zukerman M, Addie RG (1999) Application of the M/pareto process to modeling broadband traffic streams. In: IEEE international conference on networks. ICON'99 proceedings (Cat. No. PR00243), IEEE, pp 53–58
- Ni D, Leonard JD, Guin A et al (2005) Multiple imputation scheme for overcoming the missing values and variability issues in ITS data. *J Transp Eng* 131(12):931–938
- Niculescu D, Nath B (2003) Ad hoc positioning system (APS) using AOA. In: IEEE INFOCOM 2003. Twenty-second annual joint conference of the IEEE computer and communications societies (IEEE Cat. No. 03CH37428), IEEE, pp 1734–1743
- Niewiadomska-Szynkiewicz E, Sikora A, Marks M (2016) A movement-assisted deployment of collaborating autonomous sensors for indoor and outdoor environment monitoring. *Sensors* 16(9):1497
- Nurvitadhi E, Sim J, Sheffield D, et al (2016) Accelerating recurrent neural networks in analytics servers: comparison of FPGA, CPU, GPU, and ASIC. In: 2016 26th international conference on field programmable logic and applications (FPL), IEEE, pp 1–4



- Nurvithadhi E, Venkatesh G, Sim J, et al (2017) Can FPGAs beat GPUs in accelerating next-generation deep neural networks? In: Proceedings of the 2017 ACM/SIGDA international symposium on field-programmable gate arrays, ACM, pp 5–14
- Oke PR, Allen JS, Miller RN et al (2002) Assimilation of surface velocity data into a primitive equation coastal ocean model. *J Geophys Res: Oceans* 107(C9):5–25
- Oord Avd, Dieleman S, Zen H, et al (2016) Wavenet: a generative model for raw audio. arXiv preprint [arXiv:1609.03499](https://arxiv.org/abs/1609.03499)
- O’Shea TJ, Clancy TC, McGwier RW (2016) Recurrent neural radio anomaly detection. arXiv preprint [arXiv:1611.00301](https://arxiv.org/abs/1611.00301)
- Parisot S, Ktena SI, Ferrante E, et al (2017) Spectral graph convolutions for population-based disease prediction. In: International conference on medical image computing and computer-assisted intervention, Springer, pp 177–185
- Pathan ASK, Lee HW, Hong CS (2006) Security in wireless sensor networks: issues and challenges. In: 2006 8th international conference advanced communication technology, IEEE, pp 6–pp
- Picone M, Amoretti M, Zanichelli F (2012) Simulating smart cities with DEUS. In: Proceedings of the 5th international ICST conference on simulation tools and techniques, ICST (Institute for Computer Sciences, Social-Informatics and Technology), pp 172–177
- Pilosu L, Fileppo F, Scopigno R (2011) RADII: a computationally affordable method to summarize urban ray-tracing data for VANETs. In: 2011 7th international conference on wireless communications. Networking and mobile computing, IEEE, pp 1–6
- Pradhan PM, Panda G (2012) Connectivity constrained wireless sensor deployment using multiobjective evolutionary algorithms and fuzzy decision making. *Ad Hoc Netw* 10(6):1134–1145
- Proakis JG, Salehi M (2001) Digital communications, vol 4. McGraw-Hill, New York
- Pyrkov TV, Slipensky K, Barg M et al (2018) Extracting biological age from biomedical data via deep learning: too much of a good thing? *Sci Rep* 8:1–11. <https://doi.org/10.1038/s41598-018-23534-9>
- Qu L, Li L, Zhang Y et al (2009) PPCA-based missing data imputation for traffic flow volume: a systematic approach. *IEEE Trans Intell Transp Syst* 10(3):512–522
- Quan Z, Cui S, Sayed AH et al (2008) Optimal multiband joint detection for spectrum sensing in cognitive radio networks. *IEEE Trans Signal Process* 57(3):1128–1140
- Radhakrishnan N, Kandeepan S (2020) An improved initialization method for fast learning in long short-term memory-based Markovian spectrum prediction. *IEEE Trans Cogn Commun Netw* 7(3):729–738
- Radhakrishnan N, Kandeepan S, Yu X et al (2021a) Performance analysis of long short-term memory-based Markovian spectrum prediction. *IEEE Access* 9:149582–149595
- Radhakrishnan N, Kandeepan S, Yu X, et al (2021b) Soft fusion based cooperative spectrum prediction using LSTM. In: 2021 15th international conference on signal processing and communication systems (ICSPCS), IEEE, pp 1–7
- Ramos P, Santos N, Rebelo R (2015) Performance of state space and ARIMA models for consumer retail sales forecasting. *Robot Comput-Integr Manuf* 34:151–163. <https://doi.org/10.1016/j.rcim.2014.12.015>
- Rokach L, Maimon OZ (2008) Data mining with decision trees: theory and applications, vol 69. World Scientific, Singapore
- Saad W, Han Z, Poor HV et al (2012) A cooperative Bayesian nonparametric framework for primary user activity monitoring in cognitive radio networks. *IEEE J Sel Areas Commun* 30(9):1815–1822. <https://doi.org/10.1109/JSAC.2012.121027>
- Sanderson BM, Knutti R, Caldwell P (2015) Addressing interdependency in a multimodel ensemble by interpolation of model properties. *J Clim* 28(13):5150–5170
- Santana EFZ, Bastista DM, Kon F, et al (2016) SCSimulator: an open source, scalable smart city simulator. In: Tools session of the 34th Brazilian symposium on computer networks (SBRC). Salvador, Brazil, p 46
- Schmidt M, Block D, Meier U (2017) Wireless interference identification with convolutional neural networks. In: 2017 IEEE 15th international conference on industrial informatics (INDIN), IEEE, pp 180–185
- Shared Spectrum Company (2010) General survey of radio frequency bands—30 MHz to 3 GHz version 2.0. Tech. rep., shared spectrum company. [http://www.sharespectrum.com/wp-content/uploads/2010\\_0923%20General%20Band%20Survey%20-%2030MHz-to-3GHz.pdf](http://www.sharespectrum.com/wp-content/uploads/2010_0923%20General%20Band%20Survey%20-%2030MHz-to-3GHz.pdf)
- Sharma SK, Lagunas E, Chatzinotas S et al (2016) Application of compressive sensing in cognitive radio communications: a survey. *IEEE Commun Surv Tutor* 18(3):1838–1860
- Shi X, Chen Z, Wang H et al (2015) Convolutional LSTM network: a machine learning approach for precipitation nowcasting. In: Cortes C, Lawrence ND, Lee DD et al (eds) Advances in neural information processing systems, vol 28. Curran Associates Inc, Red Hook, pp 802–810



- Shi L, Zhang Y, Cheng J, et al (2018) Adaptive spectral graph convolutional networks for skeleton-based action recognition. arXiv preprint [arXiv:1805.07694](https://arxiv.org/abs/1805.07694)
- Signoretto M, Van de Plas R, De Moor B et al (2011) Tensor versus matrix completion: a comparison with application to spectral data. *IEEE Signal Process Lett* 18(7):403–406
- Sithamparanathan K, Giorgetti A (2012) Cognitive radio techniques: spectrum sensing, interference mitigation, and localization. Artech House, Norwood
- Song X, Kanasugi H, Shibasaki R (2016) DeepTransport: prediction and simulation of human mobility and transportation mode at a citywide level. In: *IJCAI*, pp 2618–2624
- Spencer P, Aultman-Hall L (2019) Assessing seasonal and climate-related variability in rates of walking and physical activity with time use data. University of Vermont Transportation Research Center, Tech. rep
- Sridhara V, Bohacek S (2007) Realistic propagation simulation of urban mesh networks. *Comput Netw* 51(12):3392–3412
- Srinivasan S, Janedula P, Dhoble S, et al (2019) High performance scalable FPGA accelerator for deep neural networks. arXiv preprint [arXiv:1908.11809](https://arxiv.org/abs/1908.11809)
- Subramaniam S, Reyes H, Kaabouch N (2015) Spectrum occupancy measurement: an autocorrelation based scanning technique using USRP. In: 2015 IEEE 16th annual wireless and microwave technology conference (WAMICON), IEEE, pp 1–5
- Sullivan O, Gershuny J (2018) Speed-up society? Evidence from the UK 2000 and 2015 time use diary surveys. *Sociology* 52(1):20–38
- Sun X (2016) Asynchronous parallel learning for neural networks and structured models with dense features. In: *Proceedings of COLING 2016, the 26th international conference on computational linguistics: technical papers*, pp 192–202
- Sun H, Laurensen DI, Wang CX (2010) Computationally tractable model of energy detection performance over slow fading channels. *IEEE Commun Lett* 14(10):924–926
- Sutton PD, Nolan KE, Doyle LE (2008) Cyclostationary signatures in practical cognitive radio applications. *IEEE J Sel Areas Commun* 26(1):13–24
- Taillandier P, Gaudou B, Grignard A et al (2019) Building, composing and experimenting complex spatial models with the GAMA platform. *Geoinformatica* 23(2):299–322
- Taylor SJ, Letham B (2018) Forecasting at scale. *Am Stat* 72(1):37–45
- Tian Z, Giannakis GB (2007) Compressed sensing for wideband cognitive radios. In: 2007 IEEE international conference on acoustics, speech and signal processing-ICASSP'07, IEEE, pp IV–1357
- Tridgell S, Boland D, Leong PH, et al (2020) Real-time automatic modulation classification using RFSoc. In: 2020 IEEE international parallel and distributed processing symposium workshops, IPDPSW 2020, New Orleans, LA, USA, May 18–22, 2020. IEEE, pp 82–89. <https://doi.org/10.1109/IPDPSW50202.2020.00021>
- Turin GL, Clapp FD, Johnston TL et al (1972) A statistical model of urban multipath propagation. *IEEE Trans Veh Technol* 21(1):1–9
- Valenzuela R (1993) A ray tracing approach to predicting indoor wireless transmission. In: *IEEE 43rd vehicular technology conference*, IEEE, pp 214–218
- Van Lint J, Hoogendoorn S, van Zuylen HJ (2002) Freeway travel time prediction with state-space neural networks: modeling state-space dynamics with recurrent neural networks. *Transp Res Rec* 1:30–39
- van der Hoog S (2017) Deep learning in (and of) agent-based models: a prospectus. [arXiv:1706.06302](https://arxiv.org/abs/1706.06302)
- Vaswani A, Shazeer N, Parmar N, et al (2017) Attention is all you need. *CoRR* abs/1706.03762. <http://arxiv.org/abs/1706.03762>
- Veen J, Van Der Wiellen P (2003) The application of matched filters to PD detection and localization. *IEEE Electr Insul Mag* 19(5):20–26
- Vieira MA, Vieira LF, Ruiz LB, et al (2003) Scheduling nodes in wireless sensor networks: a voronoi approach. In: 28th annual IEEE international conference on local computer networks, 2003. LCN'03. Proceedings., IEEE, pp 423–429
- Viriyasitavat W, Boban M, Tsai HM et al (2015) Vehicular communications: survey and challenges of channel and propagation models. *IEEE Veh Technol Mag* 10(2):55–66
- Walden MC, Rowsell FJ (2005) Urban propagation measurements and statistical path loss model at 3.5 GHz. In: 2005 IEEE antennas and propagation society international symposium, IEEE, pp 363–366
- Wang D, Yang Z (2016) An novel spectrum sensing scheme combined with machine learning. In: 2016 9th international congress on image and signal processing. BioMedical engineering and informatics (CISP-BMEI), IEEE, pp 1293–1297
- Wang Z, Salous S (2011) Spectrum occupancy statistics and time series models for cognitive radio. *J Signal Process Syst* 62(2):145–155
- Wang C, Gong L, Yu Q et al (2016) DLAU: a scalable deep learning accelerator unit on FPGA. *IEEE Trans Comput-Aided Des Integr Circuits Syst* 36(3):513–517

- Wang J, Zhou Y, Sun Y et al (2019) Cloud server oriented FPGA accelerator for long short-term memory recurrent neural networks. *J Phys: Conf Ser* 1284:012044
- Ward JA, Evans AJ, Malleson NS (2016) Dynamic calibration of agent-based models using data assimilation. *R Soc Open Sci*. <https://doi.org/10.1098/rsos.150703>
- Wellens M, Mähönen P (2010) Lessons learned from an extensive spectrum occupancy measurement campaign and a stochastic duty cycle model. *Mob Netw Appl* 15(3):461–474
- Wen Q, Gao J, Song X, et al (2019) RobustSTL: a robust seasonal-trend decomposition algorithm for long time series. In: Proceedings of the AAAI conference on artificial intelligence, pp 5409–5416
- Williams BM (2001) Multivariate vehicular traffic flow prediction: evaluation of ARIMAX modeling. *Transp Res Rec* 1776(1):194–200
- Work D, Tossavainen OP, Blandin S, et al (2008) An ensemble Kalman filtering approach to highway traffic estimation using GPS enabled mobile devices. In: 47th IEEE conference on decision and control, Cancun, Mexico, pp 5062–5068. <https://doi.org/10.1109/CDC.2008.4739016>
- Wu CH, Lee KC, Chung YC (2007) A delaunay triangulation based method for wireless sensor network deployment. *Comput Commun* 30(14–15):2744–2752
- Xie S, Liu Y, Zhang Y et al (2010) A parallel cooperative spectrum sensing in cognitive radio networks. *IEEE Trans Veh Technol* 59(8):4079–4092
- Xu K, Ba J, Kiros R, et al (2015) Show, attend and tell: neural image caption generation with visual attention. In: International conference on machine learning, pp 2048–2057
- Xue J, Feng Z, Zhang P (2013) Spectrum occupancy measurements and analysis in Beijing. *IERI Procedia* 4:295–302
- Yang H, Yang J, Han LD et al (2018) A kriging based spatiotemporal approach for traffic volume data imputation. *PLoS ONE* 13(4):e0195957
- Yasdi R (1999) Prediction of road traffic using a neural network approach. *Neural Comput Appl* 8(2):135–142
- Ye Z, EL-Azouzi R, Jimenez T, et al (2014) Computing the quality of experience in network modeled by a Markov modulated fluid model. *arXiv preprint arXiv:1412.2600*
- Yoon Y, Kim YH (2013) An efficient genetic algorithm for maximum coverage deployment in wireless sensor networks. *IEEE Trans Cybern* 43(5):1473–1483
- Younis M, Akkaya K (2008) Strategies and techniques for node placement in wireless sensor networks: a survey. *Ad Hoc Netw* 6(4):621–655
- Yu B, Yin H, Zhu Z (2017) Spatio-temporal graph convolutional networks: a deep learning framework for traffic forecasting. *arXiv preprint arXiv:1709.04875*
- Yu D, Deng L (2010) Deep learning and its applications to signal and information processing [exploratory dsp]. *IEEE Signal Process Mag* 28(1):145–154
- Yu R, Zhang Y, Yi L et al (2012) Secondary users cooperation in cognitive radio networks: balancing sensing accuracy and efficiency. *IEEE Wirel Commun* 19(2):30–37
- Zeng Y, Liang YC (2008) Spectrum-sensing algorithms for cognitive radio based on statistical covariances. *IEEE Trans Veh Technol* 58(4):1804–1815
- Zhan C, Zeng Y, Zhang R (2017) Energy-efficient data collection in UAV enabled wireless sensor network. *IEEE Wirel Commun Lett* 7(3):328–331
- Zhang S (2012) Nearest neighbor selection for iteratively kNN imputation. *J Syst Softw* 85(11):2541–2552
- Zhang W, Itoh K, Tanida J et al (1990) Parallel distributed processing model with local space-invariant interconnections and its optical architecture. *Appl Opt* 29:4790–4797. <https://doi.org/10.1364/AO.29.004790>
- Zhang C, Wu D, Sun J, et al (2016a) Energy-efficient CNN implementation on a deeply pipelined FPGA cluster. In: Proceedings of the 2016 international symposium on low power electronics and design, ACM, pp 326–331
- Zhang J, Zheng Y, Qi D, et al (2016b) DNN-based prediction model for spatio-temporal data. In: Proceedings of the 24th ACM SIGSPATIAL international conference on advances in geographic information systems, ACM, p 92
- Zhang S, Du Z, Zhang L, et al (2016c) Cambricon-X: an accelerator for sparse neural networks. In: The 49th annual IEEE/ACM international symposium on microarchitecture, IEEE Press, p 20
- Zhang H, Feng S, Liu C, et al (2019) CityFlow: a multi-agent reinforcement learning environment for large scale city traffic scenario. *arXiv:1905.05217*
- Zhao R, Song W, Zhang W, et al (2017a) Accelerating binarized convolutional neural networks with software-programmable FPGAs. In: Proceedings of the 2017 ACM/SIGDA international symposium on field-programmable gate arrays, ACM, pp 15–24
- Zhao Z, Chen W, Wu X et al (2017b) LSTM network: a deep learning approach for short-term traffic forecast. *IET Intell Transp Syst* 11(2):68–75. <https://doi.org/10.1049/iet-its.2016.0208>
- Zhou Y, Dai Z, Hao X et al (2017) Coalition formation games for cooperative spectrum sensing in cognitive radio networks. In: Handbook of cognitive radio. Springer, Singapore, pp 1–32

Zhou H, Wang H, Li X et al (2018) A survey on mobile data offloading technologies. *IEEE Access* 6:5101–5111

Zhou Y, Li T, Shi J et al (2019) A CEEMDAN and XGBOOST-based approach to forecast crude oil prices. *Complexity* 2019:1–15. <https://doi.org/10.1155/2019/4392785>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.