

Heatmap.Live: Crowd Mobility Tracking and Forecasting Using IoT Sensors and Transformers

Steven Gao hgao@exeter.edu

Phillips Exeter Academy

Abstract—Crowd counting and mobility tracking supply valuable insights into public safety, event organization, and urban planning, in a time when the COVID-19 pandemic renders the demand for a reliable and privacy-preserving crowd tracking system more exigent. Existing models struggle to capture long-range context dependency, notwithstanding a shorter forecast horizon and lower accuracy. In this paper, we develop a scalable Internet-of-Things sensor network that receives IEEE 802.11n Wi-Fi probe requests, a proven indicator of crowd activity. We then propose a data processing framework based on Google’s Transformer architecture with multi-headed attention to generate forecasts on upcoming mobility patterns, utilizing one-hot encoding and embedding layers to better express periodicity in our model. Conducting comprehensive experiments using self-collected and public data sets, our model supports a large forecast horizon and outperforms existing LSTM variants by a wide margin.

Keywords: crowd counting, mobility, IoT sensor, Wi-Fi probe request, deep learning, time series forecasting, transformer, multi-headed attention

I. INTRODUCTION

Overcrowding is a prevalent issue in a modern-day society where populations are highly concentrated – busy streets of urban areas, college campuses, shopping malls – presenting serious challenges in management and planning to administrators, and, in the meantime, compromising the convenience and safety of people traveling in or occupying such spaces. A well-designed crowd counting and mobility tracking system alleviates this pressure and provides insights on an administrative level, which could benefit public safety, event organization, urban planning, etc.

In light of the COVID-19 pandemic, the demand for a reliable crowd counting system has come into prominence when coincided with social distancing protocols that diminish the capacity of many public spaces and creates uncertainty in pedestrian traffic, businesses and schools struggle in their reopening and disease control efforts. Accurate forecasting of crowd mobility bears even more significance as administrators can take preventive measures to limit, divert, or redirect upcoming traffic flows.

Conventionally, crowd counting is achieved through video cameras and image recognition a methodology with inherent shortcomings, which is the case demonstrated in Figure 1. However, many would resent deploying security cameras on a large scale as they undermine personal privacy, and the limited range indoors mandates that more cameras are installed to cover a specific area. Environmental conditions such as low-light and, for outdoor usage, unexpected weather all impact the performance and accuracy of vision-based systems.




Fig. 1: Crowd counting with CNN (Convolved Neural Network) [1]

Image recognition and the computing resource thereof needed, combined with expensive video cameras, portray it as a costly and unsuitable solution for many use cases.

The challenge remains for constructing a reliable crowd counting system as several factors could potentially undermine the usability of such systems developed, including accuracy, coverage, cost, and privacy preservation. Whereas researchers have explored thoroughly video cameras and image recognition (e.g. [21], [11], [12], [1]), fewer have explored utilizing radio signals to determine crowd density.

The widespread and amplitude of commodity Wi-Fi devices elevates the usefulness of Wi-Fi traffic to a new level. With few exceptions where mobile devices are prohibited, almost all pedestrians now carry IEEE 802.11n Wi-Fi-capable devices such as their smartphones and tablets. Devices broadcasts probe requests to gather information on available access points nearby [5]. Probe requests could serve as an indicative metric of crowd activity at a given location.

Wi-Fi-based crowd counting addresses the limitations of vision-based systems, as the typical coverage ranges from 20 to 50 meters [14]. Wi-Fi bases systems can efficiently determine individuals’ presence through counting the MAC (media access control) addresses embedded in the probe request. To address the population without smartphones, the count can simply be extrapolated by a scalar to fit the actual circumstance. The sensors cost less to manufacture and operate in all weather and lighting conditions.

Crowd monitoring alone could render little assistance without a robust forecasting system that can inform administrators of upcoming dynamics. Previously, researchers have experimented with RNNs (Recurrent Neural Networks) [18], [13]. However, these models have inherent drawbacks – they struggle to capture long-range context dependencies, encounter issues like gradient vanishing and explosion while training, and are unable to process input in parallel, areas in which the Transformer shines.

The counter these aforementioned shortcomings and prob-

lems, we propose an interconnected network of sensors, built upon the Raspberry Pi platform. To approximate the crowd size and yield meaningful data, sensors need to be non-intrusive yet highly precise. We then develop an API server that supports distributed computing to collect and process input data from the sensors. Finally, we construct a model based on Google’s Transformer architecture, outlined in Attention is All You Need [16]. First, we test the model on data sets self-collected on the campus of Phillips Exeter Academy, and then compare our model with existing LSTM based models using a publicly available data set. We study the results of these experiments and demonstrate the efficacy of the model.

The main contributions of our work can be summarized as:

- 1) We design a scaleable and affordable IoT sensor network especially fit for usages on campuses and urban areas;
- 2) We propose a server framework that efficiently processes millions of input from the sensor network;
- 3) The Transformer-based model we constructed yields higher accuracy and better performance than existing LSTM models.

The following sections are organized as such: Section II examines the previous work in the two fields of interest – crowd counting methodologies and time series forecasting with machine learning, Section III details our design of different part of the system, Section V puts the system to test in real-world settings and examines the results in detail, Section VI concludes the paper and discusses prospective studies that can be conducted in this direction.

II. RELATED WORKS

A. Crowd Counting

1) Vision-Based Crowd Counting

Driven by big data and machine learning, a variety of crowd counting techniques have drawn attention from researchers and industry practitioners. Researchers have long delved into the possibility of using camera vision to calculate crowd counts, and while many techniques rely on face recognition of individuals and aggregating that result [14], Mikel Rodriguez et al. was among the first to propose associate density information with individual positions in the image [9]. Lokesh Boominathan, Srinivas S S Kruthiventi, and R. Venkatesh Babu employed a combination of deep and shallow CNN to capture body and facial features in combination with low-level features of blobs and clusters [1]. Although this research betters the performance of the image recognition model, other weaknesses of this method have yet to be challenged.

2) Wi-Fi Based Crowd Counting

As researchers explore other methodologies in crowd counting, Wei Xi et al. were among the first to investigate 802.11n Wi-Fi CSI (channel state information) by the rationale that CSI is highly sensitive to disruptions in the environment [20]. The investigation was furthered by Simone Di Domenico et al., who analyzed the shape of the Doppler spectrum. Han Zou et al. optimized this approach by selecting the most representative feature of CSI and applying a kernel transfer learning kernel learning to adjust for the environmental and temporal discrepancies [22]. However, two sensors are required to be

placed on two opposite ends to capture the activity in an allocated space in the approach with CSI, and experiments with groups larger than a dozen people have yet to be seen. Furthermore, CSI is highly dependent on the placement of the sensors, and change in the layout of the space could require re-fitting the data.

Wi-Fi probe requests came into spotlight with the work of J. Scheuner et al. and E. Vattapparamban et al. Probe requests are much less sensitive to alterations in the environment, expanding the coverage as a single sensor, equipped with a high-gain antenna, could capture signals within a large radius. Though probe requests cannot capture detailed activities and body gestures as CSI could, they are perfectly adequate for gathering crowd sizes. Processing probe requests is less strenuous compared to that of CSI as the unique MAC addresses can be aggregated to the sum of people.

B. Multivariate Time Series Forecasting

Many sequence DNN (Deep Neural Networks) models strive to tackle the challenge of time series forecasting, in which the input of a model is a time-dependent sequence and so is the output. Although NLP (Natural Language Processing) models enjoyed massive success in recent years, there have been incremental advancements in multivariate time series forecasting even though the two problems share many commonalities.

1) Recurrent Neural Network




Fig. 2: Typical RNN Cell [15]

RNN is apt to process sequential data like (x_1, \dots, x_n) where t is the time step index and $x_t \in R^d$. A typical RNN cell, whose structure once unrolled resembles that of lists and sequences, is shown in Figure 2. The RNN cell takes x_t as an input, and the output h_t is calculated from the input x_t and the previous hidden state. The cell runs the same function recurrently for every element in the list, allowing information from previous time steps to pass through.

2) Long-short Term Memory




Fig. 3: LSTM Cell [15]

When training recurrently over long sequences, typical RNN experiences gradient explosion and gradient vanishing, which

LSTM mediates [6]. In contrast to a typical RNN cell, an LSTM cell has three cell gates that allow it to selectively add and remove information from the cell state, as illustrated in Figure 3.

Nonetheless, while addressing some deficiencies on long-range dependencies, LSTM still has its limitations. The hidden state of a time step strongly influences only the following few time steps, resulting in a challenge to capture temporal dependencies, which is critical in time series forecasting. For instance, the crowd count from 24 hours ago might be just as indicative of as that of the last hour. Due to the nature of LSTM as a Recurrent Neural Network, data is passed in sequentially, and the next time step could only be calculated from the hidden state of the previous one, rendering parallelism an impossible task.

3) Transformer

Breaking these constraints of LSTM, Ashish Vaswani et al. proposed a new architecture in deep learning, namely Transformer, which features an encoder-decoder structure and two attention mechanisms: self-attention and encoder-decoder attention. Since it processes all time steps in parallel, Transformer outperforms LSTM in long-range context dependencies while drastically improving the efficiency.

III. SENSOR NETWORK ARCHITECTURE

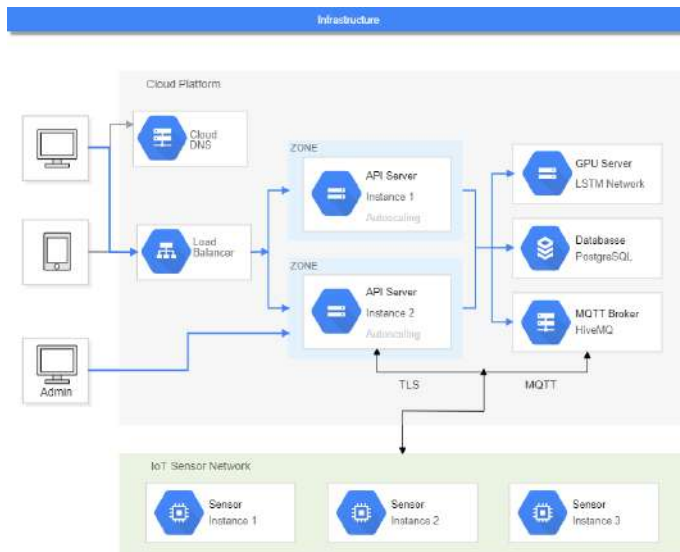


Fig. 4: System Architecture

As illustrated in Figure 4, the infrastructure of the system includes three main components: the IoT sensor network, an API server that collects and processes the data, and a Transformer-based neural network that generates forecasts of upcoming traffic dynamics.

A. Crowd Counting with IoT Sensors

1) Preliminaries: Wi-Fi Probe Request

Under the IEEE 802.11n protocol, Wi-Fi devices send out probe requests to gather information on available wireless networks nearby.[5] Probe requests are broadcast at various




Fig. 5: The Wi-Fi Probe Request Frame [5]

frequencies that are dependent on the operating system of the device. On average, smartphones send out 55 probe requests per minute and 2000 per hour in an experimental study conducted in 2015 [4]. As most people carry at least one Wi-Fi enabled device with them when traveling, probe requests become a reliable metric of crowd density.

802.11n Wi-Fi operates on 2.4GHz with 14 channels available. [5] We set the wireless module that is part of our sensor to monitor mode, and our software processes the probe requests locally. Due to limitations of available hardware, the wireless module can only listen to one channel at a time, so we iterate all of the channels on a Weighted Round Robin basis, in which process the most active channels are prioritized. We extract the MAC address and signal strength from the probe requests captured, whose structure is shown in Figure 5, and aggregated the data from our sensors to a server periodically. To exclude stationary Wi-Fi devices located near the sensor, devices, identifiable by the MAC addresses, that are always present over an extended period of time will be excluded from calculations.

2) Hardware Architecture




Fig. 6: A Sensor Node Prototype

The sensor network consists of an array of nodes, each has capabilities to monitor the traffic in the region it is installed. To ensure expandability and robustness, the sensor runs on a network operating system equipped with a run-time environment, such as Linux distributions, for our self-developed software. To monitor nearby Wi-Fi traffic, the sensor interacts with a Wi-Fi module that supports 802.11n monitor mode. [19].

We select Raspberry Pi Model 3 A+ [2] as the microprocessor of the sensor node, as it features a 64-bit SoC (System on a Chip) processor operating at 1.4GHz, 512 megabytes of RAM, and built-in Wi-Fi to report data to the server. Enclosed in a custom 3D-printed case, the Raspberry Pi communicates with an external Wi-Fi module through standard USB protocol. The Wi-Fi module is connected to a high-gain

omnidirectional antenna to increase the effective range. An optional cellular module can be connected to the Raspberry Pi to allow communication in the absence of Wi-Fi coverage.

3) Software Architecture

The sensor node software captures probe requests from the external antenna and upload the data to the server through either Wi-Fi or cellular network. Kali Linux[7] is selected as the operating system, for it provides driver support for packet capturing.

Written in Java, the sensor node software listens exclusively to probe requests and demonstrated predictable behavior and performance across different devices and environments. The sensor nodes communicate with the API server through two secured channels: MQTT over UDP and HTTP 1.2 with TLS encryption. The sensor defaults to using TLS as their means of communication with the server and falls back to MQTT, a lightweight publish-subscribe protocol suitable for transmission over less ideal network environments, reporting its operational status as well as locally hashed mac addresses.

B. Data Collection and Processing

The API server is run on a cloud server to collect and process the data it receives from the sensor network. Integrated with a web user interface, the API server is written in php using the web framework Laravel [8] to allow easy configuration for administrators. MySQL is selected as the backend to provide high-performance database service to the API server. The API server computes crowd counts and feeds the data to our neural network, detailed in Section IV, and records the data into the database. The mobile application also receives information from the API server and provides real-time information on crowdedness and forecasts to the general public.

IV. MOBILITY FORECASTING WITH TRANSFORMER-BASED MODEL

A. Data Pre-processing

Raw sensor data is processed per the following procedures.

- 1) Probe requests with a locally administered MAC address are removed as smart devices randomize their MAC addresses to protect the privacy of their users.
- 2) Devices located outside of a bounding box is eliminated. The distance of a device from the sensor is calculated using the signal strength of the probe request using Free-Space Path Loss, defined as

$$FSPL(dB) = 20\log_{10}(d) + 20\log_{10}(f) + C \quad (1)$$

where d is distance and f is frequency.

- 3) Stationary devices that are not smartphones, laptops, or other portable electronics, are filtered out based on extracted manufacturer information from the MAC address.
- 4) The raw data entries are grouped by an interval of 5 minutes and aggregated into crowd sizes based on unique MAC addresses, effectively the number of Wi-Fi devices, during that interval. As the percentage of

people who do not carry smartphones does not fluctuate frequently, the crowd count is extrapolated by a scalar to approximate the actual number of people present.

- 5) Finally, we export Data Sets 1 and 2 listed in Table I from our database as input of our model using a sliding window of a fixed look-back period and a forecast horizon.

B. Transformer-Based Model

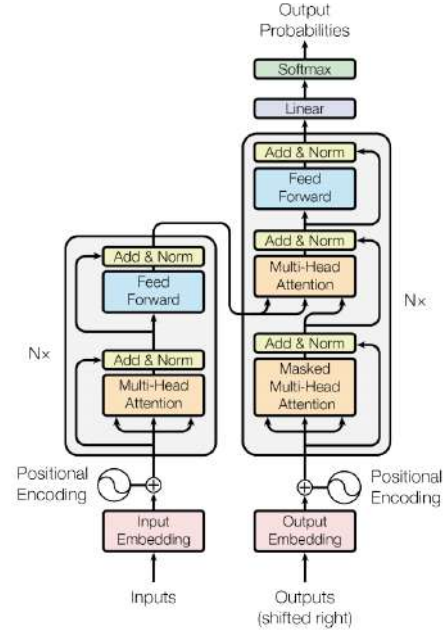


Fig. 7: The Original Transformer Architecture [16]

Illustrated in Figure 8, our Transformer-based crowd forecasting model retains the origin encoder-decoder design in Figure 7. Originally designed for Natural Language Processing, we tailor the Transformer model to fit our specific use case.

1) Encoder

The encoder accepts a sequence $(x_{t-(L-1)}, \dots, x_t)$ where $x_t \in \mathbb{R}^d$, in which d is the number of features and L represents the look-back period on historic data. The periodicity of the sequence is extracted in the One_hot encoding before the sequence is fed into a trainable Embedding Layer that formulates a representation of the time step in a global scope. Positional encoding is added for the Transformer cell to learn temporal features and dependencies within the sequence. A typical Transformer encoder block then accepts the processed sequence and passes it through multi-headed self-attention and a feed-forward network to the decoder block.

2) Decoder

Similarly, the decoding side accepts a sequence $(x_t, \dots, x_{t+(H-1)})$ where $x_t \in \mathbb{R}^d$, in which H represents the forecast horizon. The sequence serves as input for the decoder and is processed the same as it was for the encoder. The decoder has a similar design to the encoder with two attention mechanisms – self-attention and encoder-decoder attention. The self-attention mechanism works similarly to that of the encoder, and the encoder-decoder attention block

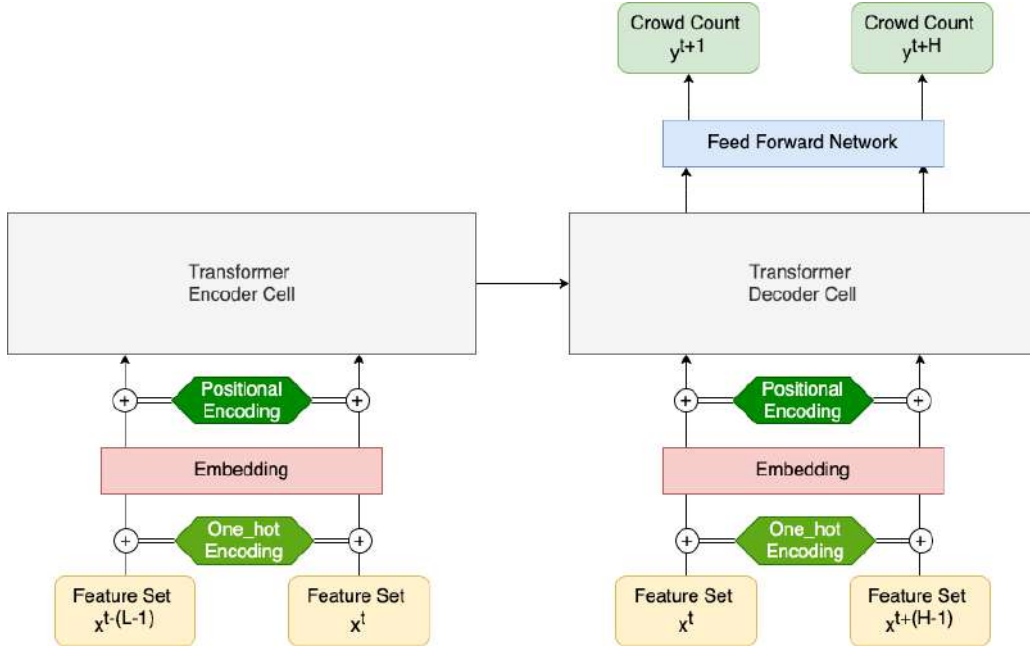


Fig. 8: Our Transformer-Based Model

takes the concatenated input from both the encoder block and the decoder input. The decoder runs on an auto-regressive basis, and, for time step t , the decoder takes the generated output of y^{t-1} as input.

C. Sequence Order and Periodicity

1) Sequence-wide Positional Encoding

The attention mechanism considers each time step as equal if without positional encoding, and the time-dependent information would be lost as a result. In Attention is All You Need, Ashish Vaswani et al. ensured that the network captures the time dependency by adding alternating sine and cosine functions to the input of both the encoder and the decoder. The period of the sine and cosine function increases from 2π to $10000\pi i$ as the dimension increases, allowing the network to study the position of a time step across the whole sequence, as shown in equations 2 and 3 [16].

$$PE_{(pos,2i)} = \sin(pos/10000^{2i/d_{model}}) \quad (2)$$

$$PE_{(pos,2i+i)} = \cos(pos/10000^{2i/d_{model}}) \quad (3)$$

2) Capturing Periodicity with One_hot Encoding

The Transformer model can keep track of long-range context dependency but is not attentive to the periodicity of the input sequence. For many cases in crowd forecasting, crowd mobility has explicit patterns in regard to the time of day, the day of the week, and whether its a weekday or a weekend. More variables, including class schedule and weather information, could be passed in as a state or condition that is present periodically. Thus, we add a layer of one_hot encoding on top of the embedding layer to extract this periodicity.

V. EXPERIMENTS AND RESULTS

In this section, we conduct experiments to test our model and evaluate its performance by comparing to LSTM models.

A. Data Sets

We evaluate our model based on four real-world data sets, including two self-collected sets on the campus of Phillips Exeter Academy and two public sets collected in Belgium in the research of Utkarsh Singh et al., as described in Table I. With maps available in Appendix I, Data Set 1 and 2 corresponds to sensor locations S1 and S4 respectively in Figure 14. Data Set 3 corresponds to sensor locations S14 to S18 in Figure 15; and Data Set 4 corresponds to sensor locations S1 to S7 in Figure 16.[13]

It is worth noting the scarcity of publicly available data sets in Wi-Fi crowd monitoring, as Data Sets 3 and 4 collected in the research of Utkarsh Singh et al. are the only ones available to the best of our knowledge.

Splitting our data into training and testing sets of different proportionality based on test scenarios, we first review Data Sets 1 and 2 in a campus setting and then Data Sets 3 and 4 in a city setting. For the evaluation of Data Sets 3 to 4, we replicate the input sequence length and horizon of [13] to ensure comparability, though our model is capable of predicting with much longer input sequences and horizons larger than that of the experiments of Utkarsh Singh et al.

B. Experimental Details

1) Training and Hyperparameters

While training with ADAM as our optimizer, teacher-forcing is utilized so that the ground truth, rather than the generated output from the last time step, is passed as input to the decoder to improve training speed. We decided against using dropout layers, as it may hinder the performance of our model on relatively small data sets. Instead, we employ kernel regularization to prevent over-fitting.

TABLE I: Data Sets Utilized in Experiment

Data Set	Ref	Location	Setting
Data Set 1	Self-collected	Downer Family Fitness Center, Phillips Exeter Academy	Campus
Data Set 2	Self-collected	Passage, Phillips Exeter Academy	Campus
Data Set 3	[13]	La Bourse de Bruxelles, Bruxelles, Belgium	Urban
Data Set 4	[13]	Sainte-Catherine, Bruxelles, Belgium	Urban

Through grid search and 5-fold cross validation, we finalized our model with 4 attention heads and 6 attention layers. Look-ahead masking is implemented in the decoder self-attention block to prevent the decoder fitting to future time steps. That is, the attention relevance scores are set to zero for future time steps, which restricts the decoder from reading future information to make predictions.

To prevent our model from making predictions solely on based on patterns manifested in the one_hot encoding, we apply to all attention layers L1 regularization, expressed as

$$L_1 = (wx + b - y)^2 + \lambda|w| \quad (4)$$

where λ is 0.0001, and L2 regularization,

$$L_2 = (wx + b - y)^2 + \lambda w^2 \quad (5)$$

where λ is 0.001, adding additional penalty to the loss function.

2) Metrics

To evaluate the performance of our crowd forecasting model, we select MSE as our loss function during training, defined as

$$MSE = \sum_{i=0}^{H-1} \frac{(\hat{y}_i - y_i)^2}{H} \quad (6)$$

where \hat{y} is the ground truth and y_i is the predicted crowd size, and evaluate our model performance using RMSE:


$$RMSE = \sqrt{\sum_{i=0}^{H-1} \frac{(\hat{y}_i - y_i)^2}{H}} \quad (7)$$

C. Forecasting with Self-Collected Data Sets

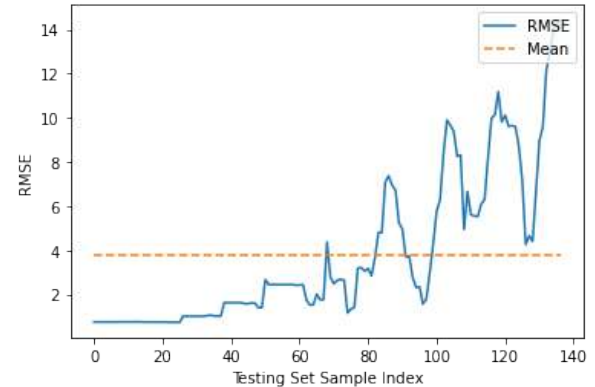
We put our model to test first in a campus setting. Using Data Sets 1 and 2, we observe whether the model could capture periodic patterns that are explicit due to a fixed class schedule and if the model is capable of discerning different patterns on weekdays and weekends. From a pragmatic standpoint, we choose 4 hours as both the look-back length and the forecast horizon as 4 hours is a sensible duration for administrators and the public to make informed decisions. For both data sets, the train test split is set to 0.8 while ensuring that the model has studied at least one day of the weekend so that it is able to recognize the different patterns. The data sets are not shuffled and are, rather, split by date.

D. Forecasting with Public Data Sets

For Data Set 3 and 4, we train the model on data of the first-three days, and test the model on five days' worth of data.



(a) Crowd Count



(b) RMSE by Testing Sample Index


Fig. 9: Data Set 1 Results

E. Results

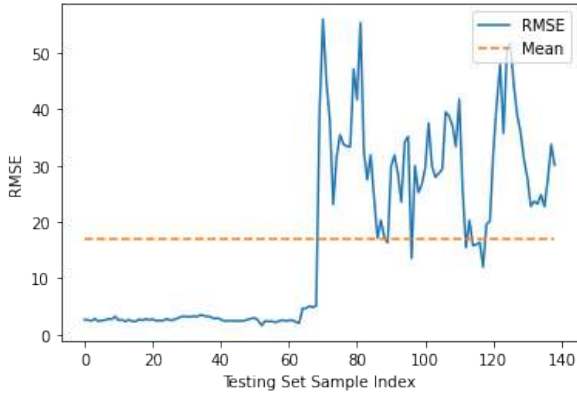
In this section, we examine the results from our own data sets, and then compare our model to existing models proposed in the recent research, Crowd forecasting based on WiFi sensors and LSTM neural networks [13].

F. Self-Collected Data Sets

Aiming to examine the accuracy of our model in a campus setting, we first look at the results of Data Sets 1 and 2. Figure




(a) Crowd Count




(b) RMSE by Testing Sample Index

Fig. 10: Data Set 2: Ground Truth and Prediction



(a) Overall RMSE by Testing Sample Index



(b) RMSE during Peak Traffic (10:30 AM to 4:30 PM)

Fig. 11: Data Set 3

TABLE II: RMSE Comparison with Existing Models in [13]

Model	Data Set 3	Data Set 4	Percentage Improvement
RW [13]	229.40	274.46	-
LSTM [13]	158.49	214.30	26.01
BiLSTM [13]	146.59	223.17	26.61
EDLSTM [13]	149.64	197.44	31.12
CNNLSTM [13]	150.06	206.39	29.26
ConvLSTM [13]	142.68	190.18	33.94
Our Model	96.93	130.24	54.91

9a and 10a represent the ground truth for the crowd count, respectively, for the two data sets. Figure 9b and 10b represent the model's RMSE across the testing data sets. A periodic pattern surfaces when we look at the similarities between each weekday. However, though not clearly discernible in 9a, the peaks do not overlap due to a different slightly schedules of every day. Manifested in Figure 10a, a drastic increase of traffic exists on specific days of the week. When evaluating the performance of the model, it is imperative to contextualize the RMSE value in the setting represented. When compared to the level of magnitude of the traffic, the RMSE is well within a reasonable range due to the uncertain nature of crowd mobility. Our model is able to maintain a relatively low RMSE over the forecast horizon of four hours, and the results demonstrate its ability to capture mobility patterns over a long time span.


1) Comparison with Existing Models

Utkarsh Singh et al. covered five different LSTM-based networks, including the Original LSTM, Bidirectional LSTM, Encoder-Decoder LSTM, Convolutional Neural Network LSTM, and Convolutional LSTM, denoted as LSTM, BiLSTM, EDLSTM, CNNLSTM, ConvLSTM respectively in Table II, and the authors included a RW (random walk) model as a baseline [13]. We referenced their data in Table II to compare with our results.


Prior research generally confirm that LSTM grapples with long-range context dependencies, and we are unaware of the performance of LSTM models over longer input sequences presented in the work of Utkarsh Singh et al., as their input is limited crowd counts in the past 60 minutes and their forecast horizon is set to 30 minutes, a relatively short period of time for some practical use cases.

To ensure comparability, we applied the exact same training data, look-back period, and forecast horizon. For Data Set 3 our model is trained on three days of data – Dec 25, Dec 26, and Dec 27 – and tested on a set of five days; similarly, for Data Set 4, our model is trained on data from Dec 01, Dec 10, and Dec 13, and evaluated on a testing set of five days as well.

As shown in Table II, our model yields the best result across



(a) Overall RMSE by Testing Sample Index



(b) RMSE during Peak Traffic (4 PM to 8 PM)

Fig. 12: Data Set 4

the board, observing the exact same training and testing setup. We significantly lowered RMSE for these two data sets, and leads the next best-performing model, namely ConvLSTM, by a 20.97% margin compared to baseline. These results further demonstrate that our model is versatile going from a campus setting with the maximum amount of traffic being around 100 people to busy urban centers with thousands of people. For forecast horizon specific comparisons, please refer to Appendices C and D.

VI. CONCLUSION

In this paper, we have developed a full-fledged crowd tracking sensor network and offered a Transformer-based approach to generating upcoming crowd mobility forecasts. We encoded periodicity into the model using one_hot to express the periodic patterns of our data. Tested on self-collected and public data sets, our model has proven its capabilities to navigate in different settings, including college campuses and busy urban centers, for it captures long-range context dependencies in historic data, increasing the forecast horizon and accuracy by a substantial margin. Comparing our model to existing models in a 2020 study [13], we lowered RMSE by 54.91% compared to the baseline model in Table II, leading the second best performer, ConvLSTM by a wide margin of 20.97%.

With our mobile application shown in Appendix A, the sensor network is now in phased deployment at Phillips Exeter Academy.

A. Future Prospects

In our future efforts, we will explore the integration of Transformer and Graph Convolutional Network that captures the crowd mobility across different locations, which we were not able to explore due to limitations of available data sets.

REFERENCES

- [1] Lokesh Boominathan, Srinivas S S Kruthiventi, and R. Venkatesh Babu. “CrowdNet: A Deep Convolutional Network for Dense Crowd Counting”. In: *Proceedings of the 24th ACM international conference on Multimedia*. 2016, pp. 640–644.
- [2] *Buy a Raspberry Pi 3 Model A+ – Raspberry Pi*. <https://www.raspberrypi.org/products/raspberry-pi-3-model-a-plus/>. (Accessed on 09/07/2020).
- [3] Simone Di Domenico et al. “Trained-once device-free crowd counting and occupancy estimation using WiFi: A Doppler spectrum based approach”. In: *2016 IEEE 12th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*. 2016, pp. 1–8.
- [4] Julien Freudiger. “How Talkative is Your Mobile Device? An Experimental Study of Wi-Fi Probe Requests”. In: *Proceedings of the 8th ACM Conference on Security Privacy in Wireless and Mobile Networks*. WiSec ’15. New York, New York: Association for Computing Machinery, 2015. ISBN: 9781450336239. DOI: 10.1145/2766498.2766517. URL: <https://doi.org/10.1145/2766498.2766517>.
- [5] Matthew S. Gast. *802.11 wireless networks the definitive guide*; OReilly, 2013.
- [6] Sepp Hochreiter and Jürgen Schmidhuber. “Long Short-Term Memory”. In: *Neural Computation* 9.8 (1997), pp. 1735–1780. DOI: 10.1162/neco.1997.9.8.1735. eprint: <https://doi.org/10.1162/neco.1997.9.8.1735>. URL: <https://doi.org/10.1162/neco.1997.9.8.1735>.
- [7] *Kali Linux — Penetration Testing and Ethical Hacking Linux Distribution*. <https://www.kali.org/>. (Accessed on 08/21/2020).
- [8] *Laravel - The PHP Framework For Web Artisans*. <https://laravel.com/>. (Accessed on 09/08/2020).
- [9] Mikel Rodriguez et al. “Density-aware person detection and tracking in crowds”. In: *2011 International Conference on Computer Vision*. 2011, pp. 2423–2430.
- [10] J. Scheuner et al. “Probr - A Generic and Passive WiFi Tracking System”. In: *2016 IEEE 41st Conference on Local Computer Networks (LCN)*. 2016, pp. 495–502.
- [11] Vishwanath A. Sindagi and Vishal M. Patel. “A survey of recent advances in CNN-based single image crowd counting and density estimation”. In: *Pattern Recognition Letters* 107 (2017), pp. 3–16.

- [12] Vishwanath A. Sindagi and Vishal M. Patel. “CNN-Based cascaded multi-task learning of high-level prior and density estimation for crowd counting”. In: *2017 14th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*. 2017, pp. 1–6.
- [13] Utkarsh Singh et al. “Crowd forecasting based on WiFi sensors and LSTM neural networks”. In: *IEEE Transactions on Instrumentation and Measurement* 69.9 (2020), pp. 6121–6131.
- [14] Utkarsh Singh et al. “Crowd Monitoring: State-of-the-Art and Future Directions”. In: *Iete Technical Review* (2020), pp. 1–17.
- [15] *Understanding LSTM Networks*. URL: <http://colah.github.io/posts/2015-08-Understanding-LSTMs/>.
- [16] Ashish Vaswani et al. “Attention is All You Need”. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 2017, pp. 5998–6008.
- [17] E. Vattapparamban et al. “Indoor occupancy tracking in smart buildings using passive sniffing of probe requests”. In: *2016 IEEE International Conference on Communications Workshops (ICC)*. 2016, pp. 38–44.
- [18] Wei Wang et al. “Occupancy prediction through Markov based feedback recurrent neural network (M-FRNN) algorithm with WiFi probe technology”. In: *Building and Environment* 138 (2018), pp. 160–170.
- [19] Wikipedia. *Monitor mode* — *Wikipedia, The Free Encyclopedia*. <http://en.wikipedia.org/w/index.php?title=Monitor\%20mode&oldid=868693168>. [Online; accessed 07-September-2020]. 2019.
- [20] Wei Xi et al. “Electronic frog eye: Counting crowd using WiFi”. In: *33rd IEEE Conference on Computer Communications, IEEE INFOCOM 2014*. 2014, pp. 361–369.
- [21] Cong Zhang et al. “Cross-scene crowd counting via deep convolutional neural networks”. In: *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2015, pp. 833–841.
- [22] Han Zou et al. “FreeCount: Device-Free Crowd Counting with Commodity WiFi”. In: *GLOBECOM 2017 - 2017 IEEE Global Communications Conference*. 2017, pp. 1–6.

APPENDIX A
GUI AND MOBILE APPLICATION

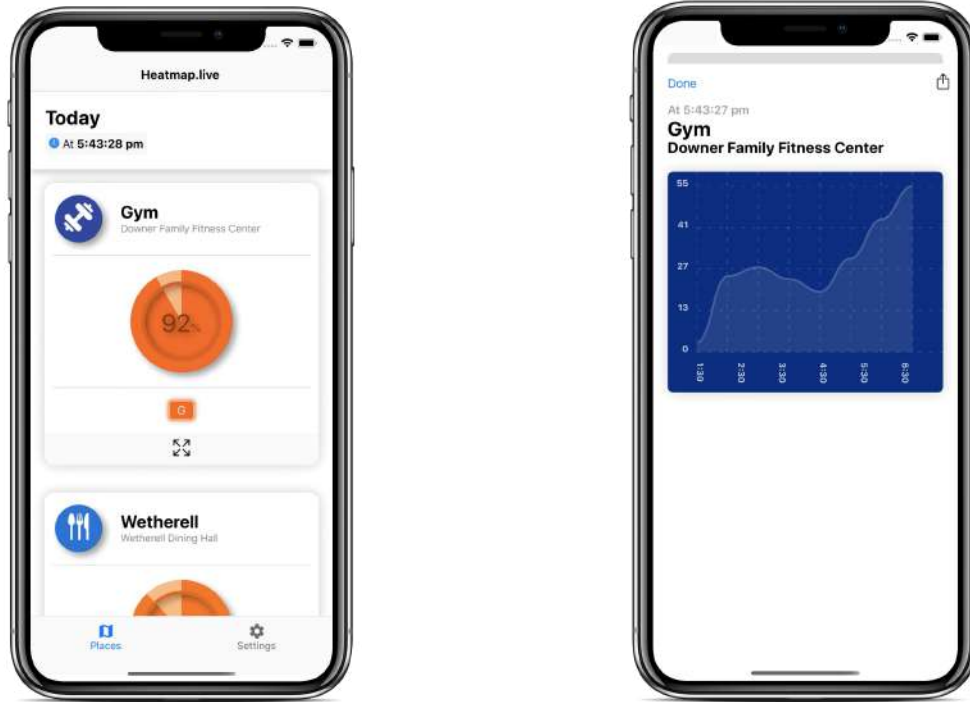


Fig. 13: Mobile Application

APPENDIX B
SENSOR LOCATIONS

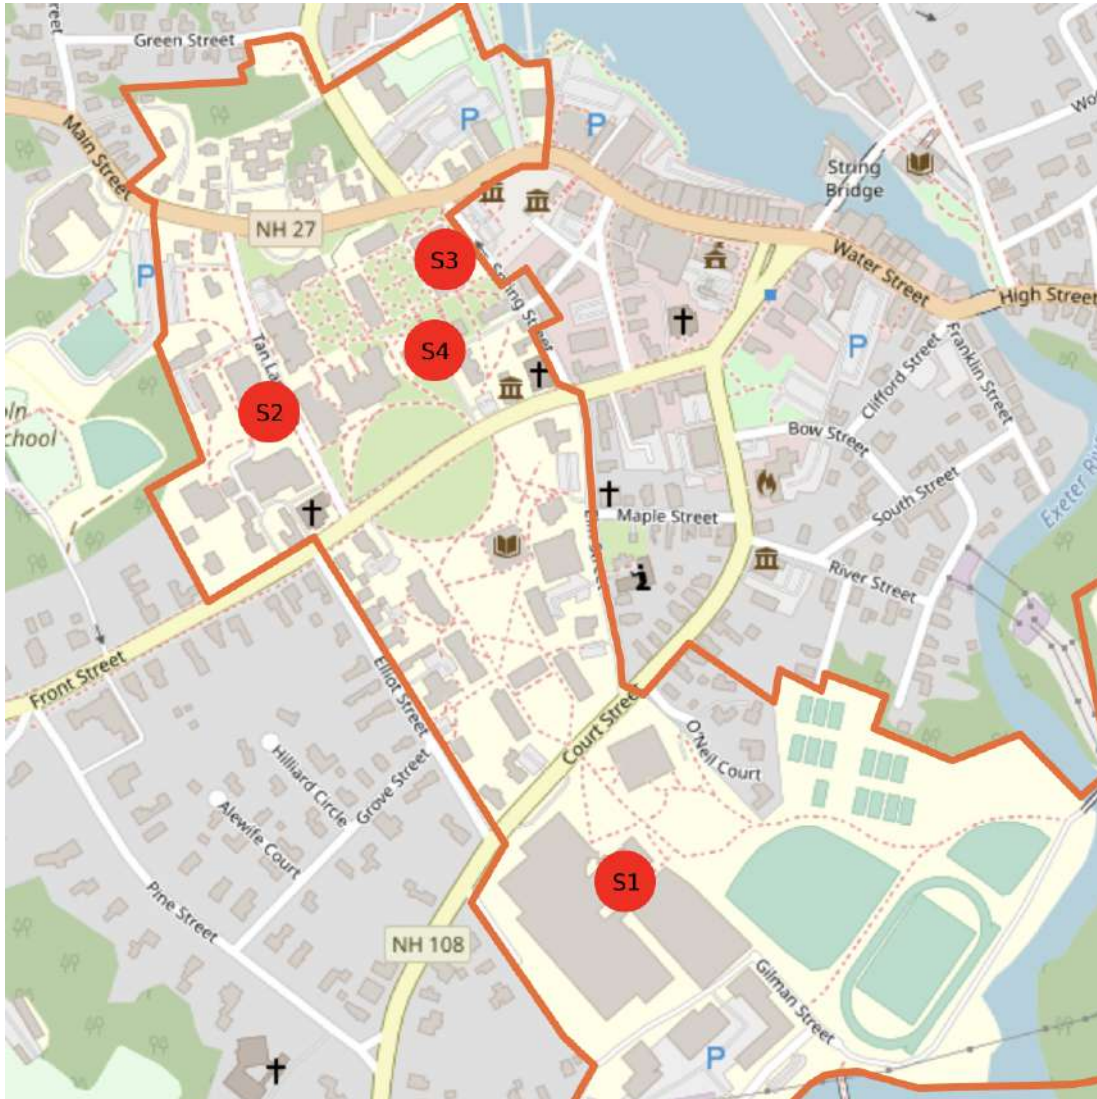


Fig. 14: Sensor Locations of Data Set 1 and 2

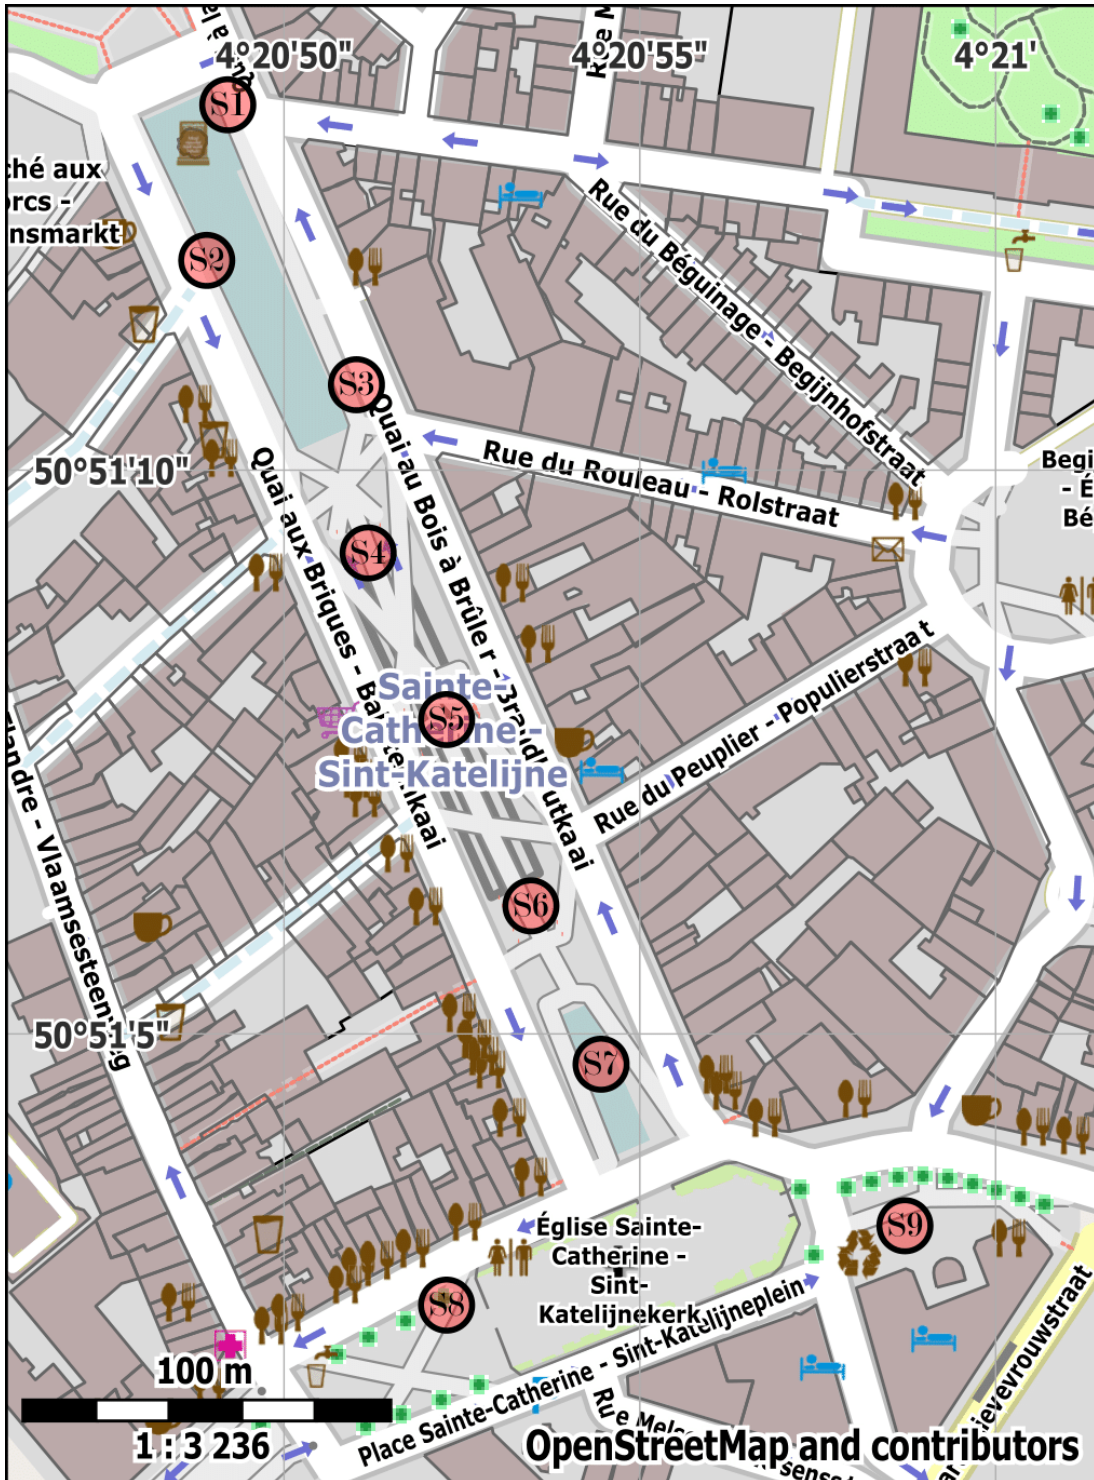



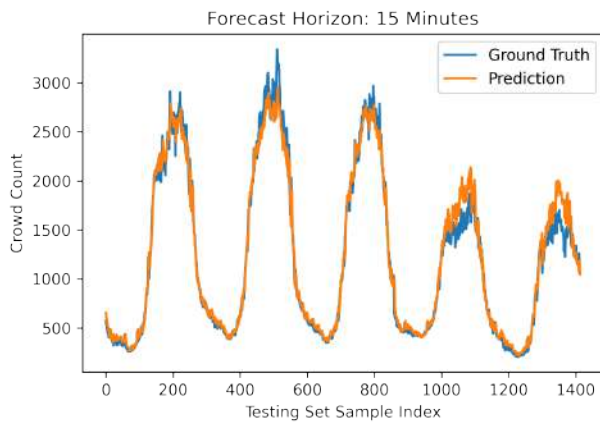



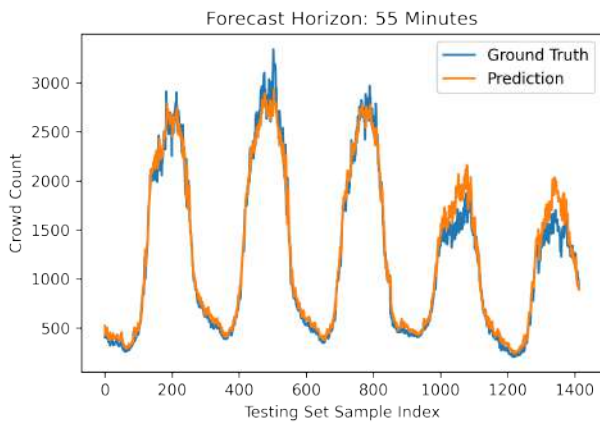


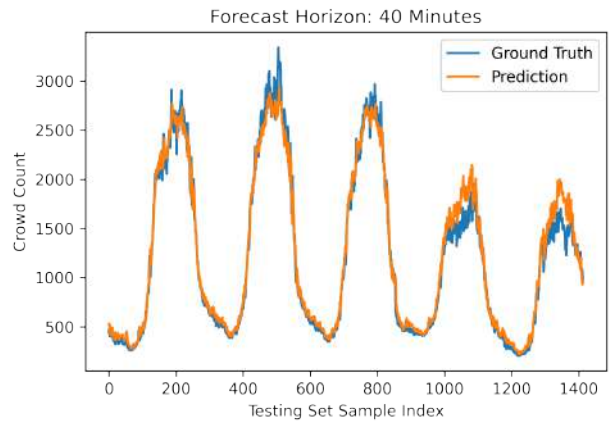



Fig. 16: Sensor locations of Data Set 4 [5]

APPENDIX C
DATA SET 3 RESULTS





APPENDIX D
DATA SET 4 RESULTS