



# Devil in the Landscapes: Inferring Epidemic Exposure Risks from Street View Imagery

Zhenyu Han \*, Yanxin Xi †, Tong Xia ‡, Yu Liu \*, Yong Li \*

\* Beijing National Research Center for Information Science and Technology (BNRist),

Department of Electronic Engineering, Tsinghua University, Beijing, China, 100084

† Department of Computer Science, University of Helsinki, Finland

‡ Department of Computer Science and Technology, University of Cambridge, Cambridge, UK

liyong07@tsinghua.edu.cn

## ABSTRACT

Built environment supports all the daily activities and shapes our health. Leveraging informative street view imagery, previous research has established the profound correlation between the built environment and chronic, non-communicable diseases; however, predicting the exposure risk of infectious diseases remains largely unexplored. The person-to-person contacts and interactions contribute to the complexity of infectious disease, which is inherently different from non-communicable diseases. Besides, the complex relationships between street view imagery and epidemic exposure also hinder accurate predictions. To address these problems, we construct a regional mobility graph informed by the gravity model, based on which we propose a transmission-aware graph neural network (GNN) to capture disease transmission patterns arising from human mobility. Experiments show that the proposed model significantly outperforms baseline models by 8.54% in weighted F1, shedding light on a low-cost, scalable approach to assess epidemic exposure risks from street view imagery.

## CCS CONCEPTS

• Applied computing → Life and medical sciences; • Computing methodologies → Computer vision; Machine learning.

## KEYWORDS

Street View Imagery, Computer Vision, Epidemic Risk, Ordinary Differential Equations, COVID-19

### ACM Reference Format:

Zhenyu Han \*, Yanxin Xi †, Tong Xia ‡, Yu Liu \*, Yong Li \*. 2023. Devil in the Landscapes: Inferring Epidemic Exposure Risks from Street View Imagery. In *The 31st ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '23), November 13–16, 2023, Hamburg, Germany*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3589132.3625596>

## 1 INTRODUCTION

With the rapid urbanization progress in the last century, more than 55% of people live in cities surrounded by the built environment

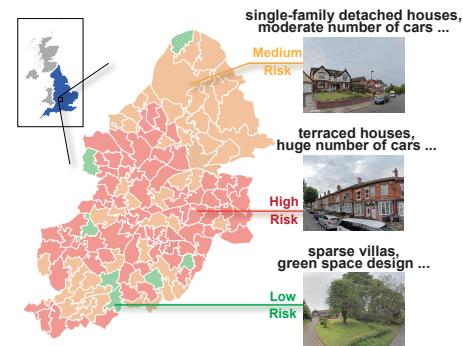


Figure 1: Illustration of epidemic exposure risk identification in Birmingham. In this study, we investigate the epidemic exposure risks for the whole England.

that provides the setting for all human activities, such as buildings, roads, green spaces, etc.. The intimacy between human and urban built environment makes it a critical environmental determinant of health [4], affecting both the physical and mental status of citizens.

With the recent development of big data processing and deep learning technology, street view imagery provides a powerful data source to assess the built environment with rich information and high scalability [7]. In Figure 1, we showcase that house types, the number of vehicles, the design of green spaces are possible transmission-related features for predicting epidemic exposure risks [9]. With publicly available street view imagery from map services or social media check-ins [8], we can predict the epidemic exposure risk in most parts of the world, even in low- and middle-income countries that may lack detailed census data.

However, accurately predicting epidemic exposure risk from street view imagery is challenging. First, infectious diseases are greatly influenced by the human mobility, which cannot be properly identified solely from street view imagery. Second, the distinct and complex transmission patterns arising from person-to-person contacts and interactions require different modeling approaches from traditional non-communicable diseases.

To overcome these challenges, we propose a novel model to identify epidemic exposure risks from street view imagery. Specifically, we construct a network of street view imagery to capture the regional transmission influence of infectious diseases, where we simulate the population flow inspired by Stouffer's law of population movement and Tobler's first law of geography. Based on



This work is licensed under a Creative Commons Attribution-ShareAlike International 4.0 License.

SIGSPATIAL '23, November 13–16, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0168-9/23/11.

<https://doi.org/10.1145/3589132.3625596>

the proposed network, we design a transmission-aware graph neural network (GNN) emulating epidemiological process inspired by Susceptible-Infectious-Recovered (SIR) model [10]. Through the proposed model, we predict the global map of epidemic exposure risk in a low-cost, scalable manner, enlightening the design methodology for creating an epidemiologically resilient living environment through the power of geographic information system (GIS).

The contributions of this study can be summarized as follows:

- We construct a regional transmission network informed by the street view imagery and human mobility, based on which the spatial correlation of epidemic can be accurately captured.
- We propose a transmission-aware GCN model with epidemiological knowledge considering the distinct transmission patterns of infectious diseases.
- We conduct extensive experiments to validate the effectiveness of the proposed model, which outperforms the best baseline by 8.54% in terms of weighted F1, 3.33% in weighted precision, and 4.93% in weighted recall.

## 2 PRELIMINARIES

The SIR model is a well-established epidemiological model that leverages the following ordinary differential equations (ODEs) to depict the dynamic of the epidemic:

$$\frac{dS(t)}{dt} = -\beta \frac{S(t)I(t)}{N}, \quad (1)$$

$$\frac{dI(t)}{dt} = \beta \frac{S(t)I(t)}{N} - \gamma I(t), \quad (2)$$

$$\frac{dR(t)}{dt} = \gamma I(t). \quad (3)$$

The above model divides the whole population  $N$  into four states:  $S$ ,  $I$ ,  $R$  for susceptible, infectious and recovered people accordingly. There are two learnable parameters:  $\beta$  is the infection rate, while  $\gamma$  is the recovery rate. It assumes a second-order contact between susceptible and infectious people for disease transmission as  $\beta S(t)I(t)$ , and a first-order natural recovery process  $\gamma I(t)$ . Leveraging the calibrated SIR model, we can estimate the basic reproduction number  $R_0$ , which reflects the transmissibility of the target disease under specific urban scenarios as follows:

$$R_0 = \beta / \gamma. \quad (4)$$

In this study, we use  $R_0$  as an agent for epidemic exposure risk.

## 3 METHODS

We illustrate the framework of this study in Figure 2. We propose a transmission-aware GCN model, *i.e.*, EpiGCN to infer the epidemic exposure risk through publicly available street view imagery. To capture the human mobility induced transmission of infectious diseases, we construct a regional mobility network, where the node feature represents geo-tagged imagery and the edge weight reflects the population flow simulated by the gravity model [11]. Inspired by the computational process of SIR ODEs, we design a novel message passing function that integrates the epidemiological model with representation learning.

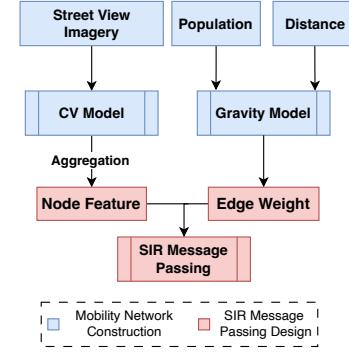


Figure 2: Framework overview.

### 3.1 Mobility Network Construction

To capture the regional contact and transmission patterns invoked by human mobility, we first construct an MSOA level mobility network as shown in Figure 3. Specifically, we have the graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where  $\mathcal{V}$  is the set of MSOA and  $\mathcal{E}$  is the set of mobility influences.

Leveraging any CV backbone model  $M_c$ , we extract the MSOA level imagery feature by aggregating all the image dense embeddings of the corresponding MSOA, which serves as the node feature  $\mathbf{h}_v, \forall v \in \mathcal{V}$ .

To depict the human mobility influence between MSOAs, we adopt the gravity model [11] to simulate regional population flows as the edge weight as follows:

$$e_{v,w} = \frac{N_v^\rho N_w^\theta}{\exp(d_{vw}/\delta)}, \forall v, w \in \mathcal{V}, \quad (5)$$

where  $N_\star$  is the population number for MSOA  $\star$ ,  $d_{vw}$  is the Euclidean distance between  $v, w$ . We set the empirical parameters  $\rho, \theta, \delta$  according to [1].

Eq.(5) predicts the population flows as proportional to neighborhood population and inversely proportional to travel distance. The design of numerator is inspired by Stouffer's law of population movement that people are more attracted by regions with more social interaction opportunities. Besides, the denominator depicts Tobler's first law of geography that people tend to visit nearby places to reduce the travel cost.

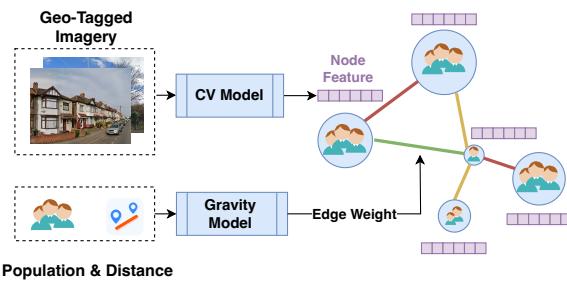
### 3.2 SIR Message Passing Design

Given the node feature  $\mathbf{h}_v$  of target node  $v$  and the neighborhood node set  $\mathcal{N}_v$ . First, we leverage three different linear layers to transform the node feature into susceptible (S), infectious (I), and recovered (R) embedding as follows:

$$S_v = \sigma(W_S \cdot \mathbf{h}_v + b_S), \quad I_v = \sigma(W_I \cdot \mathbf{h}_v + b_I), \quad R_v = \sigma(W_R \cdot \mathbf{h}_v + b_R),$$

where  $W_\star$  is the weight matrix,  $b_\star$  is the bias matrix, and  $\sigma$  is the rectified linear unit (ReLU) activation function.

The essence of SIR model lies in the second-order transmission process of  $\beta S(t)I(t)$  and the first-order recovery process of  $\gamma I(t)$ . For the transmission process, we aggregate the I embeddings from neighborhood nodes according to the edge weights generated by the



**Figure 3: Illustration of the network construction process. The size of node represents the population, and the edge color represents the strength of population flow simulated by the gravity model.**

gravity model, which will be concatenated with local  $S$  embedding to depict the second-order interaction. It goes through a linear transformation to calculate the embedding of infections. For the recovery process, the only influence factor is the local  $I$  embedding. We use another linear transformation to depict the natural recovery process that happens in infectious people. The above process is shown in the following equations:

$$S_v = S_v - W_{tran} \cdot \text{concat}\left(S_v, \sum_{w \in N_v} e_{vw} I_w\right), \quad (6)$$

$$I_v = I_v + W_{tran} \cdot \text{concat}\left(S_v, \sum_{w \in N_v} e_{vw} I_w\right) - W_{recov} I_v, \quad (7)$$

$$R_v = R_v + W_{recov} \cdot I_v, \quad (8)$$

where  $W_{tran} \in \mathbb{R}^{2D \times D}$ ,  $W_{recov} \in \mathbb{R}^{D \times D}$  represent the linear transformation for the transmission and recovery process accordingly,  $D$  is the embedding size.

To get the epidemic exposure risk prediction, we further concatenate the local  $S$ ,  $I$ ,  $R$  embedding and use another linear transformation  $W_{output} \in \mathbb{R}^{3D \times D}$  to capture the epidemic exposure risk as follows:

$$\hat{y}_v^{\text{main}} = \text{softmax}\left(W_{output} \cdot \text{concat}(S_v, I_v, R_v)\right). \quad (9)$$

## 4 EXPERIMENTS

### 4.1 Setup of the Experiment

In this work, we study 6512 middle layer super output areas (MSOAs) in England, which are fine-grained census units with a mean population of 8236 and an average area of 19.5 km<sup>2</sup>. We collect the latest street view images from Google Map for each MSOA, where we uniformly sample 9 locations within the corresponding boundary [15]. As a result, we get 215,759 images for the whole of England MSOAs, where 76.8% of them are captured after 2019. The street view images are available in 400 × 300, which will be randomly cropped into 224 × 224 before feeding into CV models. We download the MSOA level time series of COVID-19 cases in the second outbreak window (2020-09-01 to 2021-04-30) from the UK government.

In this study, we adopt the basic reproduction number  $R_0$  as an agent for epidemic exposure risk, which is a widely used metric

**Table 1: Distribution of epidemic exposure risk labels**

Category	Low Risk	Medium Risk	High Risk
Number	1727	3820	1505
Percentage	27%	52%	23%

to depict the severity of infectious diseases. Specifically, we calibrate the SIR model according to the COVID-19 time series, which provides a model-informed  $R_0$  for each MSOA. Furthermore, we categorize the extracted  $R_0$  in each MSOA according to the mean and standard deviation into three levels, which generates the epidemic exposure risk label  $r \in \{0, 1, 2\}$  for low risk, medium risk, and high risk accordingly. The distribution of labels is demonstrated in Table 1. We randomly split the dataset into training, validation, and test sets in a 6 : 2 : 2 ratio.

We implement the proposed EpiGCN in PyTorch, where we use ResNet18 initialized with ImageNet pre-trained weights as the CV backbone. Note that the whole architectures of the CV backbone are trainable. We adopt cross-entropy as the loss function. The implementation code of our model is available at <https://github.com/0showero0/EpidemicGCN>.

### 4.2 Baseline Models

To the best of our knowledge, this is the first work to identify regional epidemic exposure risks through street view imagery. We adapt three commonly used paradigms in socioeconomic prediction task to validate the proposed method: feature based baselines (BOF [12], SceneParse [6]), end-to-end supervised CV baselines (ResNet18 [5], ViT-B/32 [2]), and unsupervised baselines (Urban2vec [13], READ [3], PG-SimCLR [14]). Implementation details are summarized as below.

- **BOF** [12]: Bag of feature method leverages HOG and GIST to extract geo-tagged imagery, which follows a random forest classifier to generate predictions.
- **SceneParse** [6]: SceneParse leverages the coverage ratio of each object to train an MLP for downstream tasks.
- **ResNet18** [5]: An end-to-end deep learning CV model that follows pyramid architecture.
- **ViT-B/32** [2]: An end-to-end deep learning CV model that follows isotropic architecture.
- **Urban2vec** [13]: An unsupervised model that constructs positive and negative image pairs according to the physical distance to guide CV model learning.
- **READ** [3]: A semi-supervised model using a pretrained CV model that fine-tuned on downstream tasks using data pruning and dimensionality reduction technology.
- **PG-SimCLR** [14]: An unsupervised model that use geographical distance and POI similarity to construct positive and negative image pairs and use attention module to fuse the embeddings.

We also implement two ablation models of the proposed EpiGCN:

- **EpiGCN w/o gravity weight**: We delete the edge weight  $e_{ij}$  generated by the gravity model as described in Sec.3.1.
- **EpiGCN w/o SIR message passing**: We replace the SIR passage passing design in Sec.3.2 by vanilla GCN.

**Table 2: Performance comparison for epidemic exposure risk prediction. All the metrics are weighted ones. The average performance over 5 runs is reported. We bold the best performance, and (\*) indicates  $p < 0.01$  significance over the best baseline metrics (underlined) in ANOVA test.**

Method	F1	Precision	Recall
BOF [12]	0.4653	0.4711	0.4995
SenseParse [6]	0.4808	0.4883	0.5085
ResNet18 [5]	0.4795	0.4581	<u>0.5360</u>
ViT-B/32 [2]	0.4907	0.5195	0.5329
Urban2vec [13]	0.4870	0.4950	0.5190
READ [3]	0.4894	0.4976	0.5168
PG-SimCLR [14]	<u>0.5014</u>	0.5100	0.5219
EpiGCN w/o gravity edge weight	0.3280	0.2836	0.4268
EpiGCN w/o SIR message passing	0.4233	0.3993	0.5371
<b>EpiGCN (Ours)</b>	<b>0.5442*</b>	<b>0.5368 *</b>	<b>0.5624 *</b>

### 4.3 Performance Analysis

The overall experiment results are reported in Table 2. The baseline models are categorized into feature based baselines, end-to-end supervised CV baselines, and unsupervised CV baselines. From the results, we have the following observations and conclusions.

- Our proposed EpiGCN constantly outperforms the best baselines on all metrics statistically significantly, with 8.54% higher weighted F1 than PG-SimCLR, 3.33% higher in weighted precision than ViT-B/32, and 4.93% higher in weighted recall than ResNet18.
- Compared with the two ablation models without gravity edge weight and SIR message passing, the full model outperforms by 65.9% and 28.6% in terms of weighted F1. This phenomenon demonstrates the effectiveness of the proposed improvements.
- In general, the performance of end-to-end supervised baselines surpasses that of unsupervised baselines, which in turn outperforms feature based baselines. Supervised baselines achieve the best performance for most metrics compared with other baselines. Most of the unsupervised baselines and feature based baselines perform poorly in terms of recall metrics, which is not convincing enough for mission-critical tasks such as epidemic exposure risk identification.

## 5 CONCLUSION

In this paper, we proposed a novel model that explicitly predicts regional epidemic exposure risks through street view imagery. Considering the inherently different transmission patterns of infectious diseases, we construct a network of street view imagery linked by regional population flow, based on which we propose a transmission-aware GCN model to capture the epidemic influence arising from human mobility. In the future, we will adopt transfer learning technology to further enhance the model performance in unseen regions, which enables a low-cost, scalable approach to assess how the built environment affects disease transmission.

## ACKNOWLEDGMENTS

This work was supported in part by The National Key Research and Development Program of China under grant 2022ZD0116402, the National Nature Science Foundation of China under U22B2057, 62171260, U1936217.

## REFERENCES

- [1] Duygu Balcan, Vittoria Colizza, Bruno Gonçalves, Hao Hu, José J Ramasco, and Alessandro Vespignani. 2009. Multiscale mobility networks and the spatial spreading of infectious diseases. *Proceedings of the national academy of sciences* 106, 51 (2009), 21484–21489.
- [2] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [3] Sungwon Han, Donghyun Ahn, Hyunji Cha, Jeasurk Yang, Sungwon Park, and Meeyoung Cha. 2020. Lightweight and robust representation of economic scales from satellite imagery. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 428–436.
- [4] Zhenyu Han, Tong Xia, Yanxin Xi, and Yong Li. 2023. Healthy Cities, A comprehensive dataset for environmental determinants of health in England cities. *Scientific Data* 10, 1 (2023), 165.
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [6] Jihyeon Lee, Dylan Grosz, Burak Uzkent, Sicheng Zeng, Marshall Burke, David Lobell, and Stefano Ermon. 2021. Predicting Livelihood Indicators from Community-Generated Street-Level Imagery. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 268–276.
- [7] Tong Li, Shiduo Xin, Yanxin Xi, Sasu Tarkoma, Pan Hui, and Yong Li. 2022. Predicting Multi-level Socioeconomic Indicators from Structural Urban Imagery. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 3282–3291.
- [8] Yu Liu, Xin Zhang, Jingtao Ding, Yanxin Xi, and Yong Li. 2023. Knowledge-infused contrastive learning for urban imagery-based socioeconomic prediction. In *Proceedings of the ACM Web Conference 2023*. 4150–4160.
- [9] Quynh C Nguyen, Jessica M Keralis, Pallavi Dwivedi, Amanda E Ng, Mehran Javanmardi, Sahil Khanna, Yuru Huang, Kimberly D Brunisholz, Abhinav Kumar, and Tolga Tasdizen. 2021. Leveraging 31 million Google street view images to characterize built environments and examine County health outcomes. *Public Health Reports* 136, 2 (2021), 201–211.
- [10] Ronald Ross. 1916. An application of the theory of probabilities to the study of a priori pathometry.—Part I. *Proceedings of the Royal Society of London. Series A, Containing papers of a mathematical and physical character* 92, 638 (1916), 204–230.
- [11] Filippo Simini, Marta C González, Amos Maritan, and Albert-László Barabási. 2012. A universal model for mobility and migration patterns. *Nature* 484, 7392 (2012), 96–100.
- [12] Wenshan Wang, Su Yang, Zhiyuan He, Minjie Wang, Jiulong Zhang, and Weishan Zhang. 2018. Urban perception of commercial activeness from satellite images and streetscapes. In *Companion Proceedings of the The Web Conference 2018*. 647–654.
- [13] Zhecheng Wang, Haoyuan Li, and Ram Rajagopal. 2020. Urban2vec: Incorporating street view imagery and pois for multi-modal urban neighborhood embedding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 1013–1020.
- [14] Yanxin Xi, Tong Li, Huandong Wang, Yong Li, Sasu Tarkoma, and Pan Hui. 2022. Beyond the First Law of Geography: Learning Representations of Satellite Imagery by Leveraging Point-of-Interests. In *Proceedings of the ACM Web Conference 2022*. 3308–3316.
- [15] Guozhen Zhang, Jinhui Yi, Jian Yuan, Yong Li, and Depeng Jin. 2023. DAS: Efficient Street View Image Sampling For Urban Prediction. *ACM Transactions on Intelligent Systems and Technology* (2023).