Neural network applications to the study of air quality and health

Michelle Wan, Prof. Alex Archibald

AI4ER — UNIVERSITY OF CAMBRIDGE

Outcomes

Machine learning models were trained to predict regional daily mortality in Greater London, from input features describing air quality, meteorology and socioeconomic factors. Model performance was compared between linear and multilayer perceptron (MLP) neural network regressor models. Results found that, while both types of regressor architectures learnt to predict seasonal cycles in mortality rates, the MLP neural network made test set predictions (see right, top) with a 73% reduction in mean squared error compared to the equivalent linear model. This illustrates the improved modelling power conferred by the nonlinear nature of neural networks, despite the network here being shallow in depth.
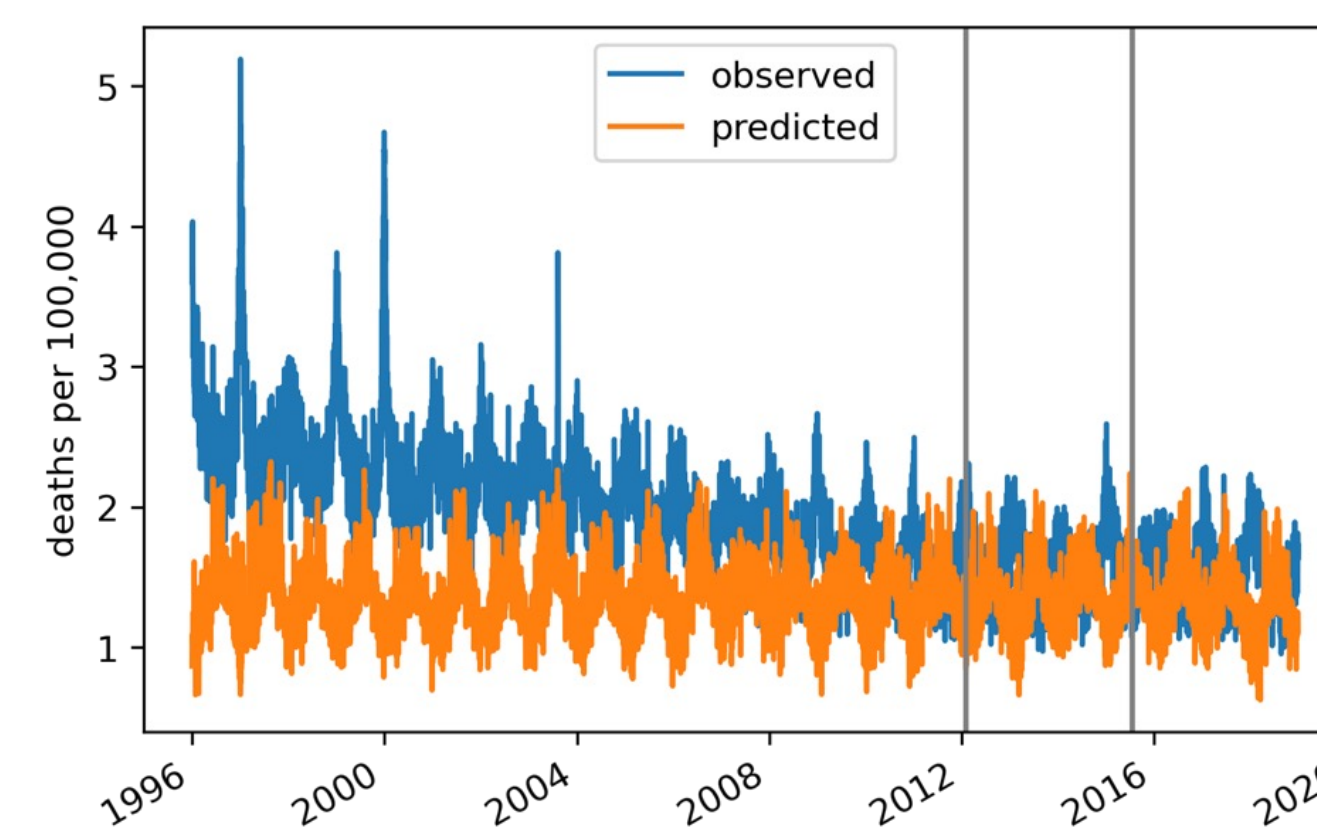
Predictive health model graphic
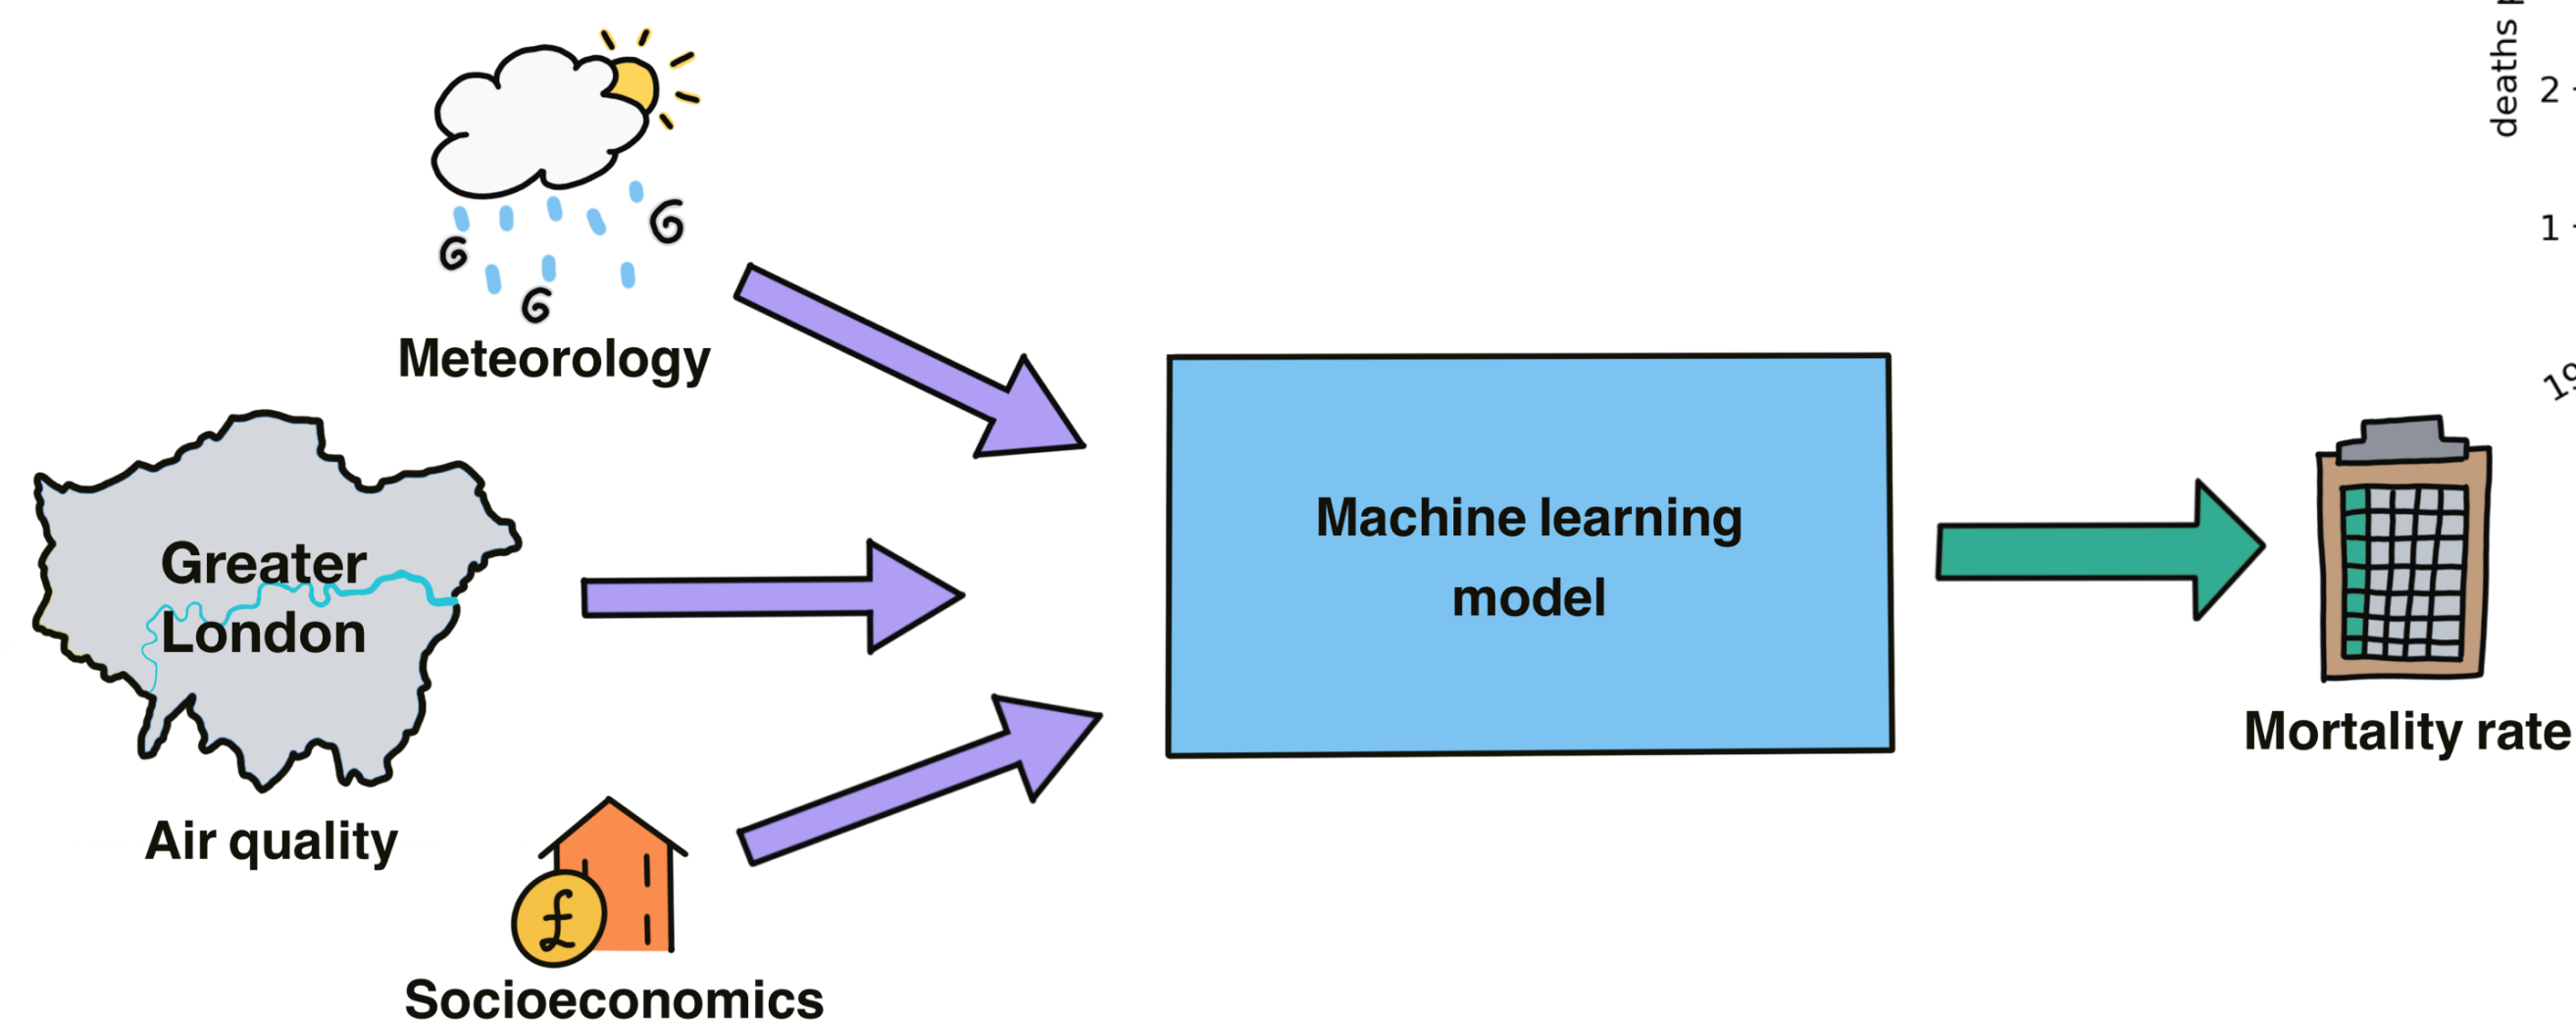
Neural network predictions

London mortality predictions by MLP regressor (all inputs)
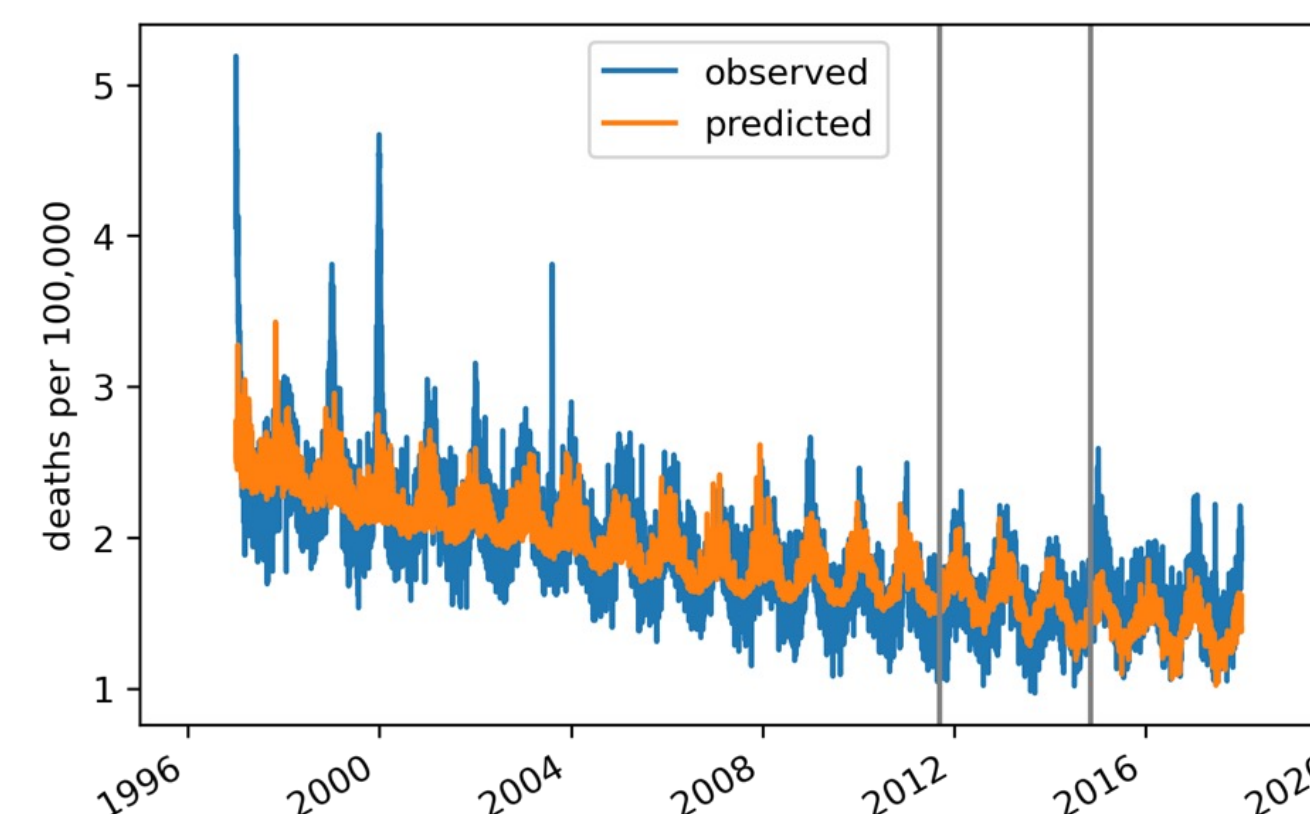
London mortality predictions by MLP regressor (excluding income)

Ablation studies

Each of three input features were excluded in turn, to test the roles they played as predictors of mortality rates in London:

– NO_2 pollutant concentration

– Temperature

– Income per capita

Results demonstrated that both types of models were dependent on the income input feature (see left, bottom) in order to accurately predict general trends in mortality rates over the two decades. Only this input feature provided information about changing trends through time, and its inclusion in this modelling approach was intended to represent the gradual improvement of societal and individual health factors.

Contact

Michelle Wan

University of Cambridge

mwlw3@cam.ac.uk