Statistical Models for Pedestrian Behaviour in Front of Bottlenecks

Nikolai Bode and Edward Codling

Abstract Understanding the movement of human crowds is important for our general understanding of collective behaviour and for applications in building design and event planning. Here, we focus on the flow of a crowd through a narrow bottleneck. We develop statistical models that describe how pedestrian behaviour immediately in front of a bottleneck affects the time lapse between consecutive pedestrians passing through the bottleneck. With this approach we isolate the most important aspects of pedestrian behaviour from a number of candidate models. We fit our models to experimental data and find that pedestrian interactions immediately in front of the bottleneck appear to be less important for the observed time lapses than interactions further away from the bottleneck. Furthermore, we demonstrate how our approach can be used to rigorously compare microscopic pedestrian behaviours across different contexts by fitting the same statistical models to three separate datasets. We suggest that our approach is a promising tool to establish similarities and differences between simulated and real pedestrian behaviour.

1 Introduction

The movement of human crowds is an important example of collective behaviour and an understanding of such systems is important for applications in building design and event planning [9]. The general consensus is that interactions between individuals are crucial to the observed dynamics at the level of the crowd. This immediately leads to two questions. First, how do individuals interact? Second, do these

Nikolai W.F. Bode
Department of Engineering Mathematics, University of Bristol, Merchant Venturers Building, Woodland Road, Bristol, BS8 1UB, UK e-mail: nikolai.bode@bristol.ac.uk

Edward A. Codling
Department of Mathematical Sciences, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK e-mail: ecodling@essex.ac.uk

interactions differ across contexts? Here, we present a framework to address both of these questions.

We focus on the paradigmatic and well-studied example of pedestrian crowds passing through narrow bottlenecks, such as exit doors. It is possible to address the aforementioned questions indirectly. For example, the average relationship between the speed of individuals and the density of crowds or the distribution of crowds in front of and inside bottlenecks provide insights into how pedestrians use the available space, whether they maintain a personal space around them and to what extent the presence of others obstructs their movement [8, 10]. Egress times can be used to investigate how different contexts (e.g. motivation of individuals) can affect the movement of a crowd [6]. Computer simulations are another well-accepted approach to explore the movement dynamics produced by different hypothesised models for interactions between individuals [9]. In contrast to such indirect approaches, our framework uses experimental, observational or simulated data to directly infer the most likely model for interactions between individuals from a set of candidate models.

For narrow bottlenecks, a conveniently measured quantity is the time lapse (or headway) between consecutive pedestrians passing through the exit. The distribution of these time lapses provides important insights: the mean is related to the pedestrian flow and the frequency of large values indicates the likelihood of jams, time-points when the flow of pedestrians slows down temporarily. Previous work has investigated properties of the distribution of time lapses. For example, it has been suggested that the distributions have a power law tail [6]. Investigating the exponent of this power law tail provides fundamental insights into the properties of the system. Exponents below a value of 2 suggest that the mean of the distribution does not converge, but grows without an upper bound as sample sizes increase - a profound issue for the prediction of bottleneck blockages. However, this scenario does not seem to apply to the pedestrian data investigated to date [6]. To study wide bottlenecks, Hoogendoorn and Daamen [8] suggest a dedicated definition of time lapses that takes the distance orthogonal to the movement direction between pedestrians into account. The authors then used a model which assumes that time lapse distributions are composed of the contribution of freely walking pedestrians and constrained pedestrians (who walk behind others) to estimate for bottlenecks of different widths the proportion of constrained and unconstrained pedestrians, as well as the average time lapse between consecutive pedestrians and the bottleneck capacity [8]. Our work presents a departure from this previous work. We develop statistical models that describe how pedestrian behaviour immediately in front of the bottleneck affects the time lapse between consecutive pedestrians passing through the bottleneck. With this framework we isolate the most likely model for pedestrian behaviour from a number of candidate models. Furthermore, we demonstrate how our approach can be used to compare microscopic pedestrian behaviours across different contexts by fitting the same statistical models to three separate datasets.

2 Methods

Let t_p and t_{p-1} be the time points at which two consecutive pedestrians cross a line that marks a physical bottleneck (e.g. an exit). We assume that the bottleneck is narrow, so that only one pedestrian at a time can exit. We define $\Delta t_p = t_p - t_{p-1}$ to be the time lapse between two consecutive pedestrians passing through this bottleneck. We propose statistical models that describe the random variable T_p which takes values Δt_p. We assume that Δt_p are distributed according to gamma distributions and the models take the general form $T_p \sim \Gamma(\mu_{p-1}, \sigma)$, where Γ denotes the gamma distribution with mean μ_{p-1} and variance σ. While we treat σ as a constant model parameter, we propose a number of models in which μ_{p-1} depends on the relative positions of pedestrians in front of the bottleneck. A different way of describing our approach is that we perform a gamma regression on the values of Δt_p using pedestrian positions in front of the bottleneck as predictors. Specifically, we use pedestrian positions at time $p-1$, i.e. at the time point when the previous pedestrian has just entered the bottleneck to predict Δt_p. The rationale behind our models is that relative pedestrian positions could affect the length of time it will take the next pedestrian to enter the bottleneck. For example, if two pedestrians are close to the bottleneck and equidistant from it, deciding who gets to exit next may take longer than if one pedestrian is much closer to the exit than the other. We propose five explanatory factors for the mean of T_p.

The first explanatory factor is a constant: $m_0 = \alpha_1$, where α_1 is a model parameter. It represents an intercept for μ_{p-1} and thus an expected baseline for Δt_p regardless of the relative positions of pedestrians.

The second explanatory factor, m_1, captures the effect on Δt_p of how densely pedestrians are clustered around the bottleneck. Let $<d>_k$ be the average distance to the exit of the k pedestrians nearest to the bottleneck (at time t_{p-1}; $k = 2,..,5$). Here and in the following, distances to the bottleneck are measured from the pedestrian position to the centre of the line pedestrians cross when entering the bottleneck. Then we define $m_1 = (\alpha_2 <d>_k - \alpha_3)^2$, where k, α_2 and α_3 are model parameters. For example, very low values of $<d>_k$ indicate that many pedestrians are very close to the bottleneck which may lead to higher expected Δt_p (pedestrians may compete against each other to exit as quickly as possible).

The third and fourth explanatory factors investigate effects of the relative positions of the two pedestrians nearest to the bottleneck. We define d_{ij} to be the difference in distance and θ_{ij} to be the angle between the two pedestrians i and j closest to the bottleneck entrance. θ_{ij} is defined as the angle between the vectors pointing from the position of pedestrians i and j to the centre of the bottleneck entrance. The third and fourth explanatory factors are defined as $m_2 = (\alpha_4 d_{ij} - \alpha_5)^2$ and $m_3 = (\alpha_6 \theta_{ij} - \alpha_7)^2$, where $\alpha_4,...,\alpha_7$ are model parameters.

The fifth explanatory factor takes the distance of the pedestrian nearest to the bottleneck, d_i, into account: $m_4 = (\alpha_8 d_i - \alpha_9)^2$ (where α_8 and α_9 are model parameters). This explanatory factor is motivated by the observation that in the absence of interactions with other pedestrians, the distance to the bottleneck of the closest pedestrian is likely to be the determining factor for Δt_p.

We use the five explanatory factors defined above to formulate sixteen models for the mean of Δt_p by considering all possible combinations of m_1, m_2, m_3 and m_4 whilst including m_0 into all models. The predicted mean for a given model is the sum over all explanatory factors included in the model. For example, the model that includes m_1 and m_3 has predicted mean $\mu_{p-1} = m_0 + m_1 + m_3$ (m_0 is included in all models). We will refer to the different models by the formula for their predicted mean (e.g. "$m_0 + m_1 + m_3$"). By comparing the extent to which these different models are supported by our data, we can establish which out of our explanatory factors best predicts Δt_p.

We use a maximum likelihood approach to fit our models. Let $f_\Gamma(\Delta t_p; \mu_{p-1}, \sigma)$ be the probability density function of a gamma distribution with mean μ_{p-1} and variance σ evaluated at Δt_p. Then we define the likelihood, L, of a given model as:

$$L = \prod_p f_\Gamma(\Delta t_p; \mu_{p-1}, \sigma), \tag{1}$$

where the product runs over all observed data points, Δt_p (i.e. time lapses). In equation 1 we assume that separate time lapses are conditionally independent given our models or that our models explain any relationships between the probability densities for separate time lapses. We describe below how we assess the validity of this assumption. For each model we find the parameter values that maximise the likelihood (requiring $\alpha_1, \alpha_2, \alpha_4, \alpha_6, \alpha_8, \sigma$ to be positive). From the maximum likelihood, L_{max}, we compute the Akaike Information Criterion (AIC), $AIC = 2n - 2log(L_{max})$, where n is the number of model parameters. We then use the AIC for model selection: models with lower AIC are better supported by the data than models with higher AIC.

We use scaled deviance residuals to assess the fit of our models to the data [5]. We plot these residuals against the predicted values, μ_{p-1}. This allows us to check if individual data points have particularly large residuals and should thus be treated as outliers. Theory predicts that the mean of the residuals should be approximately zero and systematic changes in this mean for increasing fitted values are indicative of poor model fit, suggesting that model assumptions do not hold [5]. We also plot the residuals as time series to test for correlations in our data over time that are not explained by the models. If models fit the data well, we expect that there are no systematic trends in the residuals over time.

We apply this analysis to three separate datasets. The first dataset is obtained from experiments conducted on the 1st of October 2014 in the DANA centre of the Science Museum in London with a crowd of 51 and a crowd of 71 visitors. Participants were asked to walk at a normal speed through a 0.6m wide bottleneck (1.5m long) at the end of a 2m wide corridor. At the start of the experiment, the crowd was lined up in the corridor 3m away from the bottleneck. Experiments were filmed from above at a rate of 10 frames per second and the camera (Microsoft LifeCam HD-3000 webcam^{TM}) was positioned directly above the centre of the start of the bottleneck (see figure 1a). For each frame, the position of all visible pedestrians was obtained manually by determining the position of the centre of participants' heads. Pedestrian positions were obtained in pixels and not converted to metres, as for the

purpose of our analysis only relative differences in distances matter. Time is given in seconds throughout. To reduce the sensitivity of our analysis to errors in tracked positions resulting from camera distortion, we focussed our analysis on the relative pedestrian positions directly in front of the bottleneck (e.g. low values of parameter k) where such errors are small due to the positioning of our camera. This dataset contained 325 data points (1 run with 51 participants, 4 runs with 71 participants; we only consider time lapses when at least two pedestrians are still in front of the bottleneck).

The second dataset is obtained by repeating the same experiment with a group of 39 students at the University of Bristol on the 11th of March 2015. We used the same camera, but recorded at a rate of 15 frames per second (see figure 1b for a snapshot). This dataset contained 147 data points (4 runs with 39 participants; in one trial, the exit time of the last participant was not recorded).

The third dataset is obtained from simulations of a previously published model for pedestrian movement [1]. Briefly, pedestrians move in continuous two-dimensional space. The environment (e.g. walls), as well as preferred movement directions, are encoded in a discrete floor field [4] and interactions between pedestrians and the environment (e.g. avoiding walls) are modelled via forces acting on point masses [7]. We simulate crowds of 80 pedestrians, each with $0.5m$ diameter, exiting a room through one $0.8m$ wide door (see figure 1c). At the start of simulations, pedestrians are distributed randomly inside the room. All other model details and parameters are given in previous work [1]. Simulation parameters are not fitted to the experimental data. In our simulated dataset, we aggregate the data from the first $200s$ for 10 separate simulations. This results in 717 data points.

3 Results and Discussion

Figures 1d-f show the time lapse distributions for the three datasets we analyse. We begin by discussing the results for the first dataset to show how our framework can be used to identify the most likely model for pedestrian behaviour from our candidate models. Figure 1g shows the AICs for each model obtained from our maximum likelihood fits. We can see that most models are better supported by the data (i.e. have lower AICs) than the baseline model, m_0, which assumes time lapses are independent of the relative pedestrian positions in front of the bottleneck. It is also clear that our analysis does not simply favour more complex models. For example, the model with the most parameters, $m_0 + m_1 + m_2 + m_3 + m_4$, does not have the lowest AIC. Based on our results, the model that is best supported by the data is $m_0 + m_4$, in which time lapses depend on the distance of the pedestrian nearest to the bottleneck. At first glance this appears to be a trivial result: the closer individuals are to the exit, the less time it will take them to reach it. However, this model outperforms other models for interactions immediately in front of the bottleneck implying that such interactions are less predictive for the observed time lapses. This suggests that in this data the order in which pedestrians exit is already determined when individ-
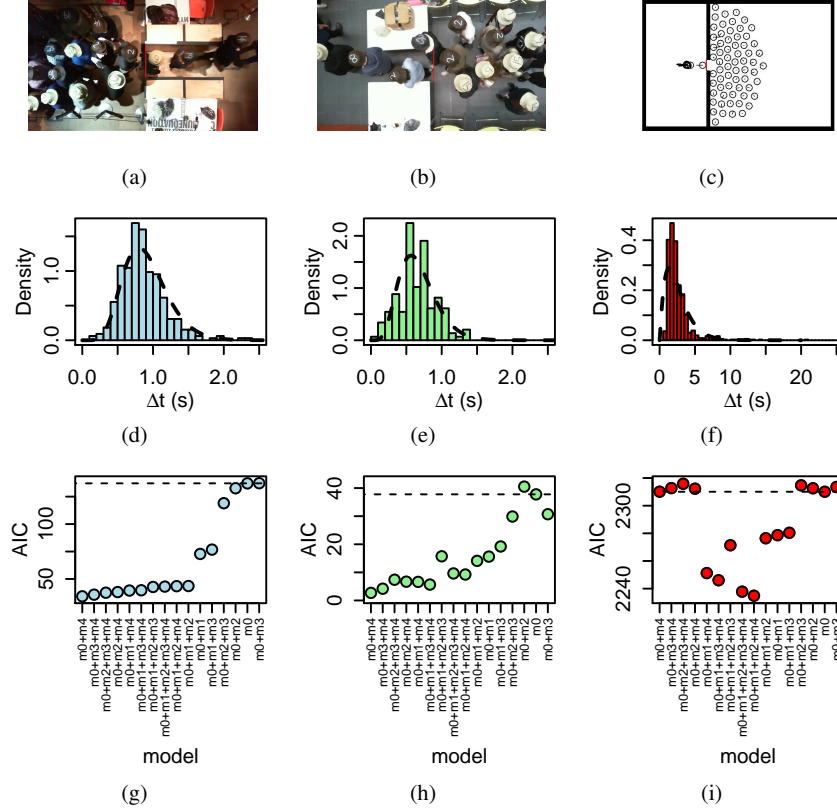
Fig. 1 Still images of experimental setup in London (a), Bristol (b) and model simulations (c). Red line segments indicate the bottleneck location used in the analysis. (d-f) time lapse distributions for the datasets presented in still images above. Dashed lines show gamma distribution fit corresponding to model m_0. (g-i) AICs for all models analysed for the three datasets. Models are shown in the same order in all panels. Dashed horizontal lines indicate the AIC of the baseline model (m_0)

uals get close to the bottleneck meaning that interactions between pedestrians occur further away from the bottleneck. We caution that while this framework establishes the relative performance of different models, it should not be over-interpreted as determining the actual, true mechanism underlying pedestrian interactions: our approach is inherently probabilistic and only investigates a potentially incomplete set of candidate models.

Importantly, our analysis allows us to compare the relative performance of the different models across datasets (figure 1g-i). General trends in AICs indicate similarities or differences across contexts. The trends in AICs for our two experimental datasets are very similar (figure 1g,h), while the trend for the simulated dataset differs substantially from the others (figure 1i). For the simulated dataset, only models that include the contributing factor m_1 perform much better than the baseline

model and the AIC of the model that is best supported in both experimental datasets, $m_0 + m_4$, does not even outperform the baseline model. The performance of models including m_1 for the simulated data suggests that the density of pedestrians in front of the bottleneck helps considerably to predict the time lapse between consecutive pedestrians. At this point we have to re-iterate that we have not fitted our simulations to the experimental data. Therefore this result should not be interpreted to mean that our model or any similar model is based on fundamentally different interaction mechanisms than seen in experimental data. Rather, we show that our framework can be used to compare microscopic pedestrian behaviours across different contexts and datasets.

An additional advantage of our framework is that residuals provide a well-established approach to assess the fit of models to data. In figure 2 we show residual plots for the models with the lowest AIC for each dataset. There does not appear to be a systematic trend in residual means against predicted values (figure 2a-c). On close inspection, we find a consistent increase in residuals over time for each simulation run in the model data (figure 2f; multiple simulation runs are shown). In the experimental data, no such temporal correlations are immediately obvious (figure 2d,e). This suggests that our model does not capture aspects of the simulated data. One explanation for this result could be that our simulations implement forces between pedestrians: as the number of pedestrians in front of the bottleneck decreases, the pressure exerted on pedestrians by others behind them is decreased which results in longer time lapses. In this way the analysis of residuals does not only highlight potential shortcomings of our statistical models, but it can also provide an additional way to highlight differences in interaction mechanisms between datasets.

4 Conclusions and Outlook

We have presented a framework of statistical models to analyse microscopic pedestrian behaviour in front of narrow bottlenecks. This approach allows us to isolate the most likely interaction mechanism from a list of possible mechanisms which can also be used to compare the behaviours underlying different simulations and datasets in a rigorous and quantitative way.

Our work presented here is a starting point and many extensions of our models to take different scenarios and behaviours into account are possible or even necessary depending on the context. For example, interactions within social groups, such as friends, could be considered [2]. Our framework could also be developed further to permit analysis of wider bottlenecks (an important topic, as wider bottleneck occur frequently and pedestrians appear to preferentially exit through wider bottlenecks [3]). Our models focus on interactions in front of bottlenecks but dynamics inside bottlenecks could have a strong influence on time lapses between consecutive pedestrians (e.g. consider a jammed bottleneck). If such dynamics are important, it is likely that our approach as presented here is not appropriate.
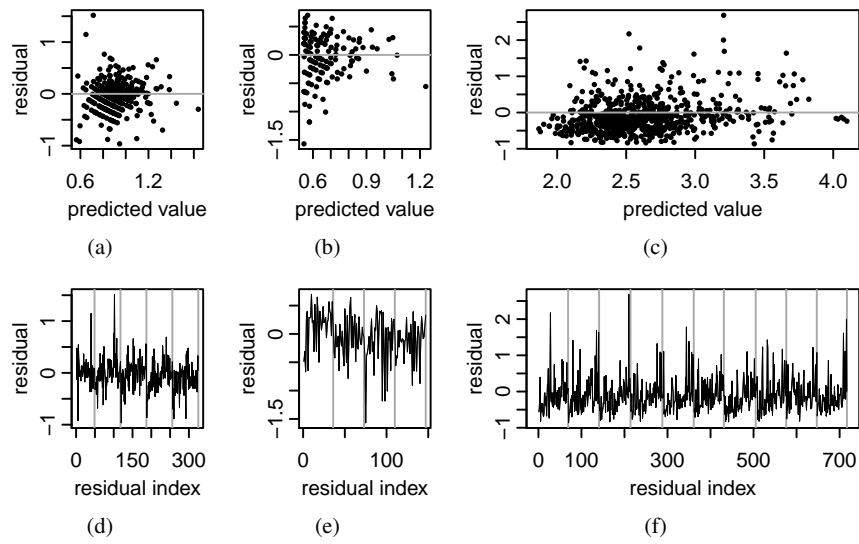
Fig. 2 Deviance residuals plotted against predicted mean time lapse for the model with the lowest AIC for the London (a), Bristol (b) and simulated (c) data. Plots of residuals ordered in time for the same data (d-f). Vertical grey lines separate data from different experimental or simulation runs

To test the usefulness of our framework, it will be necessary to apply our analysis to a wide variety of simulated, experimental and observational data. In particular, applying our analysis to experimental data in which the behaviour of individuals is controlled to some extent (e.g. motivation levels) would be very informative.

In conclusion, we suggest that our approach is a promising first step to directly establish mechanistic similarities and differences between simulated and real microscopic pedestrian behaviour.

References

1. Bode, N., Codling, E.: Human exit route choice in virtual crowd evacuations. Animal Behaviour **86**(2), 347–358 (2013)
2. Bode, N., Holl, S., Mehner, W., Seyfried, A.: Disentangling the impact of social groups on response times and movement dynamics in evacuations. PLoS ONE **10**(3), e0121,227 (2015)
3. Bode, N., Kemloh Wagoum, A., Codling, E.: Information use by humans during dynamic route choice in virtual crowd evacuations. Royal Society Open Science **2**, 140,410 (2015)
4. Burstedde, C., Klauck, K., Schadschneider, A., Zittartz, J.: Simulation of pedestrian dynamics using a two-dimensional cellular automaton. Physica A: Statistical Mechanics and its Applications **295**(3), 507–525 (2001)

5. Davison, A., Snell, E.: Residuals and diagnostics. In: D. Hinkley, N. Reid, S. EJ (eds.) Statistical Theory and Modelling. Chapman and Hall, London (1991)
6. Garcimartín, A., Zuriguel, I., Pastor, J., Martín-Gómez, C., Parisi, D.: Experimental evidence of the faster is slower effect. Transportation Research Procedia **2**, 760–767 (2014)
7. Helbing, D., Farkas, I., Vicsek, T.: Simulating dynamical features of escape panic. Nature **407**(6803), 487–490 (2000)
8. Hoogendoorn, S., Daamen, W.: Pedestrian behavior at bottlenecks. Transportation science **39**(2), 147–159 (2005)
9. Schadschneider, A., Klingsch, W., Klüpfel, H., Kretz, T., Rogsch, C., Seyfried, A.: Evacuation dynamics: Empirical results, modeling and applications. In: Encyclopedia of complexity and systems science, pp. 3142–3176. Springer (2009)
10. Zhang, J., Seyfried, A.: Quantification of bottleneck effects for different types of facilities. Transportation Research Procedia **2**, 51–59 (2014)