

Northumbria Research Link

Citation: Amos, Martyn and Webster, Jamie (2022) Crowd-Sourced Identification of Characteristics of Collective Human Motion. *Artificial Life*, 28 (4). pp. 401-422. ISSN 1064-5462

Published by: The MIT Press

URL: https://doi.org/10.1162/artl_a_00381 <https://doi.org/10.1162/artl_a_00381>

This version was downloaded from Northumbria Research Link:
<https://nrl.northumbria.ac.uk/id/eprint/48972/>

Northumbria University has developed Northumbria Research Link (NRL) to enable users to access the University's research output. Copyright © and moral rights for items on NRL are retained by the individual author(s) and/or other copyright owners. Single copies of full items can be reproduced, displayed or performed, and given to third parties in any format or medium for personal research or study, educational, or not-for-profit purposes without prior permission or charge, provided the authors, title and full bibliographic details are given, as well as a hyperlink and/or URL to the original metadata page. The content must not be changed in any way. Full items must not be sold commercially in any format or medium without formal permission of the copyright holder. The full policy is available online: <http://nrl.northumbria.ac.uk/policies.html>

This document may differ from the final, published version of the research and has been made available online in accordance with publisher policies. To read and/or cite from the published version of the research, please visit the publisher's website (a subscription may be required.)

Crowd-Sourced Identification of Characteristics of Collective Human Motion

Martyn Amos ¹ (@martyn_amos), Jamie Webster ¹,

Corresponding: Martyn Amos (martyn.amos@northumbria.ac.uk)

1. Department of Computer and Information Sciences, Northumbria University, Newcastle upon Tyne, United Kingdom.

Abstract. Crowd simulations are used extensively to study the dynamics of human collectives. Such studies are underpinned by specific movement models, which encode rules and assumptions about how people navigate a space and handle interactions with others. These models often give rise to macroscopic simulated crowd behaviours that are statistically valid, but which lack the noisy microscopic behaviours that are the signature of believable real crowds. In this paper, we use an existing “Turing test” for crowds to identify realistic features of real crowds that are generally omitted from simulation models. Our previous study using this test established that untrained individuals have difficulty in classifying movies of crowds as real or simulated, and that such people often have an idealised view of how crowds move. In this follow-up study (with new participants) we perform a second trial, which now includes a training phase (showing participants movies of real crowds). We find that classification performance significantly improves after training, confirming the *existence* of features that allow participants to identify real crowds. High-performing individuals are able to *identify* the features of real crowds that should be incorporated into future simulations if they are to be considered realistic.

Keywords: crowds, simulation, realism, agents, Turing test

1 Introduction

A significant amount of artificial life research is concerned with studying the collective dynamics of *mobile agents* operating in a spatially-explicit environment. Relevant domains include the flocking behaviour of birds and other “animats” (“boids” being the archetypal example (Reynolds, 1987)), the power of distributed swarm robotics (Brambilla et al., 2013), and the engineering of biological cell populations (Gorochoowski, 2016). In all such cases, agents (whether simulated or physically realised) are situated in Cartesian space, and may interact both with one another and with their environment.

One specific area of growing interest is the study of *crowd dynamics* (Adrian et al., 2019); that is, the behaviour of large numbers of human individuals moving through and interacting in a given environment. The need to understand collective human behaviour in physical space is pressing, as it has significant implications for events planning and management (Crociani et al., 2016), urban design (Feng et al., 2016), and incident response and analysis (Harding et al., 2011; Pretorius et al., 2015). During and after the COVID pandemic, with potentially long-lasting and profound structural and behavioural changes being made, the need to understand the crowd will persist (Pouw et al., 2020).

Due to the inherent difficulty of performing large-scale experiments with human participants, *crowd simulations* (Thalmann & Musse, 2013) (usually using an agent-based approach) are often used to investigate collective behaviour and the impact of physical or behavioural interventions on crowd dynamics. Two features of simulations are of interest; *validity* and *believability*. *Validity* describes how closely the output of the model matches data obtained from the real world (Klöpfel, 2007; Pettré et al., 2009; Seer et al., 2014). *Believability* is subtly different, and concerns the human perception of whether or not a crowd’s behaviour is *realistic*, or plausible. We are not concerned with “cinematic”, photo-realistic believability of the rendering of a crowd, but whether or not observers are able to detect characteristic *patterns of behaviour* in real crowds which are absent in simulated

27 crowds. Fundamentally, we assume that a simulation is valid, and are interested in whether
28 or not it also *looks realistic*.

29 The rest of the paper is organised as follows; we give some background motivation, outline
30 our hypothesis, and describe our crowd Turing test framework for its investigation. We
31 then describe our experimental method for the current study, and describe our results. We
32 conclude with a discussion of the implications of our findings, and suggest possible future
33 work.

34 **2 Background and Motivation**

35 Crowd simulations are now used extensively in a wide range of application domains, from
36 urban planning (Aschwanden et al., 2011), emergency response (Mahmood et al., 2017),
37 games and training simulations (Mckenzie et al., 2008), and the CGI generation of Holly-
38 wood movie scenes (a classic example being the large-scale battle scenes in *The Lord*
39 *of the Rings* series) (Ricks, 2013). Most crowd simulations are underpinned by a be-
40 havioural/movement model, which makes simplifying assumptions about individuals, and
41 which is used by agents to determine their trajectories through the simulated space.

42 The Social Forces Model (SFM) (Helbing & Molnar, 1995) lies at the heart of many scientific
43 and commercial crowd simulation packages, such as FDS+EVAC (Korhonen et al., 2010),
44 PedSim (Gloor, 2016), SimWalk (Kimura et al., 2003) and MassMotion (Rivers et al., 2014).
45 However, there are well-established deficiencies in this and other existing movement mod-
46 els. As (Lerner et al., 2007) argue, “While such approaches may capture the broad overall
47 behaviour of the crowd, they often miss the subtle details displayed by the individuals. The
48 range of individual behaviours that may be observed in a real crowd is typically too com-
49 plex for a simple behavioural model... Simple things such as walking in pairs, stopping to
50 talk to someone, changing one’s mind and heading off in a different direction or aimlessly
51 wandering about, are just a few examples which are difficult to capture.” The emphasis

52 here is less on the locomotion model of avatars or the cosmetic appearance of the agents,
53 and more on the *patterns* and “quirks” of movement that distinguish a real crowd from a
54 simulated one.

55 Why is this important? After all, emergency planners (to take one significant user group)
56 will generally be satisfied if the overall outcome of a simulation (in terms of the time re-
57 quired to evacuate a stadium, for example) is broadly valid, and will usually not concern
58 themselves with micro-level “turbulence” and other localised phenomena. However, as
59 (Fuchsberger et al., 2017) argue, crowd simulations still meet with resistance from deci-
60 sion makers in some significant industrial and societal domains, and this may be due to a
61 lack of trust in their outputs (caused, in turn, by a lack of realism). Specific concerns iden-
62 tified of relevance to the current paper include “unnatural motion paths”, so if we can go
63 some way towards addressing this, then it may lead to increased acceptance and uptake
64 of these techniques.

65 As we argue in (Webster & Amos, 2020), there is still a need for more realistic behavioural/
66 movement models in crowd simulation, and “This is motivated by a widely-acknowledged
67 need for crowd simulations to include more realistic features derived from individual and
68 social psychology (such as group-level behaviours, indecision, etc.) (Lemercier & Auberlet,
69 2016; Seitz et al., 2017; Templeton et al., 2015), which are generally not included in software
70 packages, and which give rise to rather unrealistic or “robotic” patterns of behaviour at the
71 population level”.

72 Much work has already been done on making crowd simulations more realistic; here we
73 highlight some representative contributions. (Lerner et al., 2007) describes the construc-
74 tion of a database of behavioural “motifs” which may be incorporated into an agent’s be-
75 haviour. (Peters & Ennis, 2009) used manual annotation of observations to extract in-
76 formation about group-level behaviours that were then incorporated into simulations (this
77 study also included human trials of perception of realism). More recently, (Wei et al., 2018;

78 Yao et al., 2020) used machine learning to extract features of observed crowds, which were
79 then incorporated into a crowd simulation, but neither study assessed whether or not these
80 modifications actually made the overall crowd behaviour more realistic.

81 Fundamentally, what passes for realistic is inherently subjective. To our knowledge, until
82 we performed this study no extensive work had been done on capturing the “essence” of
83 what makes a crowd realistic *from the perspective of human observers*.

84 Our previous work (Webster & Amos, 2020) showed that crowd simulations that employ the
85 most commonly-used movement model are valid (in terms of their outputs having the same
86 statistical properties as observed crowds), but they still possess a “signature” that allows
87 them to be distinguished from real crowds. Simply put, to human observers, simulated
88 crowds are still perceived differently to real crowds. Importantly, though, we also found that
89 although people are able to reliably *partition* crowds into real/simulated, *they are unable*
90 *to tell which is which*. That is, individuals are able to separate crowd movies into two
91 categories, but they are unable to reliably label the real crowds. We found that individuals
92 tend to have an idealised view of the behaviour of real crowds, which is often at odds
93 with reality. These findings confirm the observation that real and simulated crowds have
94 different microscopic features that allow them to be partitioned, if not classified.

95 To summarise, our previous work established the *existence* of features that are present in
96 real crowds but not in simulated crowds; the aim of the current paper is to *identify* those
97 features. In (Webster & Amos, 2020) we argue that “Our results suggest a possible frame-
98 work for establishing a minimal set of collective behaviours that should be integrated into
99 the next generation of crowd simulation models.” Here, we use the “Turing test” classi-
100 fication task to identify that specific set of features that allow trained viewers to reliably
101 *classify* (not just partition) real and simulated crowds. Our results show that classification
102 performance over a population of observers increases significantly after an initial training
103 phase, and that individuals are able to identify a core set of realistic behaviours that are

104 present in real crowds, but which are absent in simulated crowds. This immediately sug-
105 gests new features that must be incorporated into future crowd simulations if they are to
106 be considered realistic.

107 **3 Hypothesis**

108 In a landmark paper (Turing, 1950), Alan Turing proposed a method to investigate what
109 would become known as “artificial intelligence”. Rather than directly answering the some-
110 what ambiguous question “Can machines think?”, Turing preferred to reframe the issue in
111 terms of an “imitation game”, in which an interrogator engaged in conversation with two
112 agents via “teletypes”. One of the agents (A) is a man, and the other (B) a woman, and
113 the interrogator’s objective is to decide which is which by asking questions of both and
114 assessing their responses. The task of A is to cause the interrogator to guess *incorrectly*
115 (that is, persuade them that he is a woman), and the task of B is to “help” the interrogator
116 to guess correctly, generally by giving truthful answers. We may, therefore, interpret the
117 imitation game (commonly referred to as the “Turing test”) more generally, with the role of
118 A being played by an artificial system that seeks to persuade a human observer that it is
119 the “genuine article”, and B being played by an actual “real world” example of the system
120 under study. Importantly, the test does not seek to establish the “truth” of A’s outputs
121 (that is, their validity), but simply whether or not A could be said to represent a reasonable
122 facsimile of the system represented by B.

123 This conceptual framework has been proposed for biological modelling (Harel, 2005) and
124 artificial life (Cronin et al., 2006) as a way of investigating the realistic properties of arti-
125 ficial systems. We previously used the same approach to investigate crowd simulations,
126 basing our approach on a related Turing test for collective motion in fish (Herbert-Read
127 et al., 2015). In (Webster & Amos, 2020), we describe the results of initial experiments,
128 using a total of 540 in-person participants. The first set of trials presented individuals with

129 a sequence of paired movies, using a side-by-side representation. In each pair, one of the
130 movies represented the movement of a real crowd, and the other represented a computer
131 simulation of the same scenario (the ordering was randomised). All observations were of
132 the same physical space, and both movies were generated using the same custom render-
133 ing engine. For each pair (over six pairs in total), participants were asked to specify which
134 of the pair they thought was the real crowd (that is, they had to *identify* the real crowd). For
135 the second set of trials, participants were presented with the movies individually, and this
136 time they were asked to *classify* each movie as either real or simulated.

137 We found that participants performed better when they were asked to *classify* crowds rather
138 than having to choose between the two, but a striking feature of our results was that neither
139 mode allowed participants to perform better than random guessing. A simplistic interpre-
140 tation of this result could be that existing simulations are good enough to “pass” the crowd
141 Turing test, as human observers are unable to distinguish between them, but here we em-
142 phasise that the imitation game, as originally described by Turing, requires the interrogator
143 to be able to specify *which* agent is the man.

144 Strikingly, the most common score in the first trial was zero, meaning that a significant
145 proportion of participants (36.46%) failed to identify a single real crowd. That is, their
146 entire perception of what constitutes a real crowd was perfectly “flipped” compared to re-
147 ality. This sizeable group of participants were able to perfectly partition movies into real or
148 simulated, but were utterly unable to say which was which. This confirmed the existence of
149 a set of real crowd behaviours (informally described by participants in terms of “standing
150 around” and “moving with purpose”) that allowed individuals to separate real from sim-
151 ulated, but which were incorrectly ascribed to the simulation as generating “unrealistic”
152 crowd behaviour. Our conclusion was that participants had an idealised view of real crowd
153 behaviour, and preferred to think that it was much less “messy” and unpredictable than
154 observations would suggest.

155 Our hypothesis, therefore, is that participants in a crowd Turing test will improve their
156 classification performance after being trained by viewing real crowds, as a result of being
157 able to identify and ascribe *only to real crowds* the realistic features that are manifested
158 in the training set.

159 **4 Experimental Methods**

160 Our protocol was largely modelled on that of (Webster & Amos, 2020), but limitations
161 imposed by the COVID pandemic required us to perform our trials online, as opposed
162 to in-person. We do not believe that this modification had any significant impact on our
163 results; indeed, it actually allowed us to recruit a more diverse range of participants, rather
164 than using only University students (which was a possible criticism of the original study).

165 We performed two sets of Turing test experiments; the first (Test 1) was an online-only
166 repetition of the second (classification) test from (Webster & Amos, 2020), with entirely
167 new participants. We attracted 232 participants, who were recruited via social media.
168 This first test allowed us to assess the ability of each untrained participant to classify
169 crowds as either real or simulated, thus assigning each one a baseline score. We allowed
170 an appropriate period of time to pass (4 months) in order to ensure that the tests were
171 independent (that is, any learning effects from the first test would not be carried over to
172 the second). We then contacted every Test 1 participant who supplied an email address
173 to invite them to participate in the follow-up Test 2 (they were each offered a 10 GBP gift
174 card as an incentive); 50 participants accepted our invitation. Test 2 participants were then
175 “trained” by asking them to first watch six rendered movies of crowds that were explicitly
176 described as real. Participants then performed a second version of the classification task
177 (as in Test 1), using a different set of real and simulated clips to those used previously (in
178 order to avoid effects induced by familiarity with the clips).

179 Given that each participant had a known baseline score from Test 1, we were able to es-



Figure 1: Single movie frame of the Edinburgh Informatics Forum, taken from (Majecka, 2009).

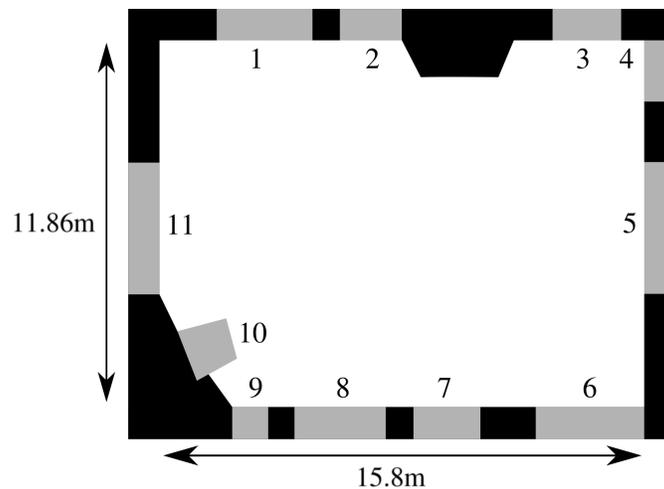


Figure 2: Diagram of Edinburgh Informatics Forum (ingress and egress points numbered), taken from (Webster & Amos, 2020)

180 tablish whether or not the training phase had a significant effect on classification ability.
 181 Participants were specifically asked to identify features that they thought allowed them to
 182 distinguish between real and simulated crowds.
 183 Test 1 was performed at the end of June-start of July 2020, and Test 2 was performed in
 184 December 2020. Our trial protocol was approved by the Northumbria University Faculty

185 of Engineering and Environment Ethics Committee, application number 24623. We now
186 describe each component of the trial in more detail.

187 **4.1 Pedestrian motion dataset**

188 As we employed the same dataset used in our previous study, we take our description of
189 it from (Webster & Amos, 2020). We used data on real pedestrians from the University of
190 Edinburgh School of Informatics (Majecka, 2009). This public dataset, captured in 2010,
191 contains over 299,000 individual trajectories corresponding to the movement of individ-
192 uals through the School Forum, and is one of the largest open datasets of its type. It has
193 been used in several studies of pedestrian movement and tracking; (Fernando et al., 2018)
194 used the dataset to pre-train short and long term trajectory prediction models, proposing
195 a “light-weight” sequential Generative Adversarial Network (GAN) architecture for person
196 localisation, which “overcomes issues related to occlusions and noisy detections”. In a
197 case study on the Edinburgh Informatics forum, (Lovreglio et al., 2017) developed a “mi-
198 croscopic calibration procedure” for floor field cellular automaton models, comparing two
199 floor field specifications to identify the best model for simulating pedestrians in the forum.
200 However, this study was only concerned with individual trajectories, and did not consider
201 the crowds as a collective. Finally, recurring activity patterns that “appear, peak, wane and
202 disappear over time” were identified using non-parametric Bayesian methods which cou-
203 ple spatial and temporal patterns with “minimal” prior knowledge (H. Wang & O’Sullivan,
204 2016).

205 **4.1.1 Environment**

206 A photo of the Forum space is shown in Figure 1, and a diagram is shown in Figure 2. The
207 Forum is rectangular in shape (measuring approximately 15.8×11.86 metres), has eleven
208 ingress/egress points, and is generally clear of obstructions. Images were captured (9
209 per second) by a camera suspended 23m above the Forum floor, from which individual

210 trajectories were extracted and made available (extraction was performed by the author of
211 (Majecka, 2009)). We note that only the *trajectories* have been made publically available,
212 and not the original video recordings, for ethical and practical reasons (these files require
213 several terabytes of storage). Importantly, none of the individuals whose trajectories were
214 captured were actively participating in movement studies; the trajectories, therefore, are
215 as close to “natural” as possible (i.e., they have “behavioural ecological validity” (Lovreglio
216 et al., 2017)).

217 **4.1.2 Pedestrian dataset**

218 The dataset is stored across a number of files, each file representing a day’s worth of crowd
219 recordings. Each file stores a list of “sightings” over that period, where a sighting is defined
220 as an individual entering (but not necessarily leaving) the frame (of course, individuals
221 may also leave and then re-enter the frame, which would be interpreted as an entirely new
222 sighting). Each row in the file therefore corresponds to a “sighting”. Every sighting during
223 the time period covered by the file is assigned a unique “agent ID”, and the individual’s
224 trajectory is stored as a list of 3-tuples of the form $\langle x, y, timestep \rangle$. Each time step codes
225 for one *frame* in the original footage (recorded at 9fps). (Majecka, 2009) note that “the
226 sample rate can vary over short periods” due to errors with the capture program; however
227 “since each captured frame is relatively independent of captured frames more than 10-20
228 seconds later”, this did not significantly impact on the quality of the resulting trajectories.

229 In what follows, we use the term “clip” to specifically refer to a time-limited sequence of
230 trajectory data (whether taken from the Edinburgh dataset or from the output of a simula-
231 tion), as opposed to a movie visualisation. We first wrote a script to convert a list of tra-
232 jectories into a frame-by-frame representation of agent locations over time. This outputs
233 co-ordinates for *all* of the visible agents at *each* time step, which is required for rendering
234 the trajectories into videos, as well as for analysing the crowds at each point in time. We
235 also wrote another script to essentially reverse this process (extracting individual trajec-

236 tories from time step data), which is necessary for analysing certain features of individual
237 trajectories in clips (both real and simulated).

238 **4.1.3 Data cleaning**

239 Occasionally lossy detection by the camera means that some trajectories have missing
240 sections for several time steps; once rendered, these individuals temporarily disappear
241 from the frame and then reappear. To address this, we automatically detected such sit-
242 uations and interpolated co-ordinates for the missing time steps when parsing the Edin-
243 burgh dataset. Each new co-ordinate is placed proportionally between the surrounding
244 co-ordinates, depending on the number of missing time steps. As the Edinburgh data
245 trajectories were recorded at 9 frames per second these additional co-ordinates prevent
246 agents from disappearing in renders, but do not alter the overall shape of trajectories.
247 Across the estimated 7.9 million coordinates in the dataset, a total of 230,046 trajec-
248 tory time gaps were identified. Of these, 128,660 (55.93%) were made up of 1 frame and
249 49,794 (21.65%) were 2 frames in duration. The largest observed time gaps were 13 and
250 14 frames; however these were each only identified once, and were not present in the real
251 crowd data clips used in this research. Approximately 99.20% of all identified time gaps
252 were of 9 frames or fewer (approximately one second of camera tracking), and interpolation
253 of these time gaps did not result in any observable issues. We also increased the number
254 of frames per second of both sets of trajectories (real and simulated), from 9 to 72, by
255 interpolating co-ordinates. This improved the “smoothness” of the trajectories once ab-
256 stracted and rendered into video clips. This enables smooth video playback for the purpose
257 of comparisons, but does not alter the shape of the trajectories, as the distance between
258 co-ordinates is negligible. Figure 3 shows all co-ordinate trajectories in one crowd clip
259 rendered to single images at both 9 and 72 frames per second.

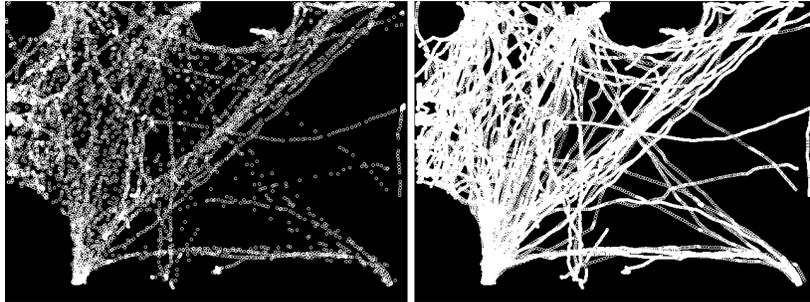


Figure 3: All trajectories in a crowd clip rendered to single images at 9 (left) and 72 (right) frames per second.

260 **4.1.4 Visualisation**

261 We wrote a utility to search the Edinburgh dataset and extract clips of a specific duration
262 containing a specific number of individuals. Both simulated and real individuals were ren-
263 dered in a uniform fashion, using a tool coded in Java. This allowed us to produce “top
264 down” visualisations of both real and simulated clips that were identical in appearance,
265 with individuals represented as filled circles, and headings depicted by an arrow (see Fig-
266 ure 4). Stationary agents in real crowd clips appear to “flick” their headings rapidly due to
267 inaccurate camera detection, so headings are only rendered when an agent is in motion.

268 The use of abstract, simplified shapes, and a top-down, two-dimensional presentation is
269 relatively common in crowd studies (N. Bode et al., 2015; Singh et al., 2009; Smith et al.,
270 2009; Wagner & Agrawal, 2014; W. L. Wang et al., 2017; Zhang et al., 2019), although three-
271 dimensional representations are also used (Loscos et al., 2003; Luo et al., 2008; Moussaïd
272 et al., 2016; Pelechano et al., 2007).

273 As in (Webster & Amos, 2020), we decided against using realistic body shape rendering
274 and 3D views, as initial tests suggested that such a presentation scheme (using animated



Figure 4: Example rendering of a crowd scene, taken from (Webster & Amos, 2020).

275 avatars) would actually distract viewers from the main aim of the experiment, which was
276 to look for *patterns of behaviour* in the crowd. Additionally, at least one study has shown
277 that crowds that are viewed from the top-down are perceived as being just as realistic as
278 those viewed from eye-level (Ennis et al., 2011).

279 The simulated crowd trajectories were converted into the same format as the real crowds
280 for rendering. Each time step has a corresponding set of co-ordinates representing a real
281 or simulated person in the Edinburgh forum, as well as their heading. At every time step in
282 a clip, our rendering tool generates a PNG image, and the sequence was then combined
283 into a video. The staircase represented in blue is an obstacle which simulated agents
284 avoid, and the staircase represented in black is an egress point located slightly inside the
285 forum.

286 **4.1.5 Clip analysis**

287 For each clip, we extracted the route choice distribution and the entry time distribution
288 for all individuals. This allowed us to initialise our simulations with the same distributions,
289 ensuring that the runs closely matched the macroscopic properties of the real-world ob-
290 servations (while leaving room for the microscopic differences in which we are interested).
291 In a later Section, we show heatmaps of the entry and exit distributions of the real crowd

292 clips.

293 After rendering real crowd clips from the Edinburgh dataset for the first time, we saw a clear
294 difference in the maximum velocity and acceleration of agents in several clips, with some
295 agents moving unnaturally quickly. This was attributed to the variability in camera capture
296 rate discussed earlier. To adjust for this variability, we calculated the average velocity of
297 individuals in each clip, and used this to scale the clip’s length (by modifying the video
298 playback speed), thus normalizing the velocity of individuals relative to expected walking
299 speed (Bohannon, 1997).

300 **4.2 Simulation construction**

301 Each test required participants to classify a number of clips of pedestrian movement as ei-
302 ther real or simulated. We began by selecting, at random, a number of clips (30s duration)
303 from the Edinburgh dataset, and extracting information about the number of individuals
304 visible and the entry/exit point distribution. This information was then used to “seed”
305 a simulation. In this way, we obtained both real and simulated versions of the same sce-
306 nario; the real version was a rendered version of the actual observations, and the simulated
307 version was a rendered version of the output of the model.

308 In order to model the scenarios captured in each real Edinburgh clip, we simulated pedes-
309 trian movement using the Vadere package (Zönnchen et al., 2020). This is an open-source
310 package, which means that (unlike commercial software) its movement models are open
311 to inspection. Importantly, it also allows for easy exporting of simulated pedestrian trajec-
312 tories, which is necessary for rendering.

313 A crucial component of the simulation is the *crowd motion model*. This defines the rules
314 of interaction between individuals (e.g., avoidance), and between individuals and their en-
315 vironment (e.g., repulsion from walls and physical obstacles), as well as route choice be-
316 haviour and differential walking speed. Many different crowd motion models exist (Duives

et al., 2013), but perhaps the most commonly-used type is based on *social forces*. Helbing and Molnar’s social force model (SFM) (Helbing & Molnar, 1995) is a microscopic, continuous model which uses “attractive” and “repulsive” force fields between individuals (and between individuals and their environment) to guide movement.

We selected the SFM as the baseline model for our simulations, as (1) it is very well-established and available for use in most open-source crowd simulation software, (2) “optimal” parameters have been refined over time, and (3) it is “recommended for pedestrian crowd movement research” following the thorough review by (Duives et al., 2013). We also compared the SFM with the Gradient Navigation Model (GNM) (Dutra et al., 2017), in order to avoid potential bias imposed by only using one motion model. The GNM is available as a default model type in Vadere, and we found that GNM simulation outputs have similar statistical properties to SFM outputs.

Table 1: Vadere simulation model parameters for SFM/GNM.

Parameter	SFM Value	GNM Value
ODE Solver	Dormand-Prince	Dormand-Prince
Pedestrian body potential	2.72	2.72
Pedestrian recognition distance	0.3	0.8
Obstacle body potential	20.1	20.1
Obstacle repulsion strength	0.25	0.25
Pedestrian radius (m)	0.2	0.2
Pedestrian speed distribution mean (m/s)	1.4	1.4
Pedestrian minimum speed (m/s)	0.4	0.4
Pedestrian maximum speed (m/s)	3.2	3.2
Pedestrian acceleration (m/s)	2	2
Pedestrian search radius (m)	1	1

For all simulations, we use the pre-supplied Vadere templates for the SFM/GNM, with default attributes and parameters (listed in Table 1). We note that all default parameter values are the same across both models, with the exception of “Pedestrian recognition distance” (0.3 for SFM, and 0.8 for GNM), but we do not believe this had any significant impact on

333 our results.

334 Vadere stores its simulation input files in JSON format, and these files specify the topog-
335 raphy of the simulation space and initial spawn parameters for each agent (or group of
336 agents). This makes it possible to write a script which generates a JSON file for each sim-
337 ulation, including the Edinburgh forum topography, as well as a JSON object for each agent
338 to be simulated. We ran each simulation in Vadere using the new simulation input files,
339 and then imported each resulting file of crowd trajectories into MATLAB to be processed.

340 In Test 1 we used only the SFM movement model; in Test 2, we divided the simulations
341 between the SFM and the GNM, in order to test whether different movement models have
342 unique movement “signatures”.

343 As discussed in (Webster & Amos, 2020), we added small amounts of noise to the sim-
344 ulated trajectories in order to replicate noise in the real crowd data. Typically, in crowd
345 videos, shoulder “swaying” can account for perceived side-to-side movement of pedes-
346 trians; however, the Edinburgh individuals were detected by an overhead camera running
347 at 9fps (placed too high to detect shoulder sway). However, occasionally faulty detection
348 caused very short-term errors in the extracted trajectories. Once rendered, this caused
349 individuals to appear to rapidly “flick” between two headings. As we had no reliable way
350 to quantify the (by inspection, small) amount of noise in the trajectories, we adjusted this
351 by eye until the apparent noise in the simulated data matched the noise level observed
352 in the real data. At any time step, a simulated agent has a 15% chance of temporarily
353 “flicking” their heading by a randomly selected value up to 45 degrees (without changing
354 their trajectory). The inclusion of noise in simulations has been shown to replicate real be-
355 haviour in animal models (N. W. Bode et al., 2010) whilst “preserving emergent behaviours
356 of previous models”. In this case, the noise added to simulated trajectories only served to
357 replicate faulty detection artefacts in the data, without altering the overall trajectories of
358 the agents.

4.3 Simulation validation

It is important to ensure that simulations (regardless of the movement model) produce outputs that are valid, so we first calculated several statistical properties for a set of simulations and the Edinburgh observations on which they were based.

As in (Webster & Amos, 2020), we used two metrics (Herbert-Read et al., 2015); *polarization* and *nearest neighbour distance* (NND). The first metric is particularly useful for describing the existence of large groups who might be moving together along the same heading (e.g., leaving a lecture room and moving together towards an exit), while the second metric is used for estimating overall crowd density. Although these metrics have tended to be used in “swarming” models (e.g., of birds or fish) in which agents are supplied with local information about other agents in their vicinity, they have recently also been used effectively to assess a model of collective behaviour based purely on vision, which is perhaps better aligned to our current model (Bastien & Romanczuk, 2020).

Polarisation measures the level of “order” in a crowd, in terms of the heading alignment of members. Polarisation is zero when the crowd is completely disordered (everyone is pointing in a different direction), and has a maximum value of 1 when all members of the crowd have the same heading:

$$\varphi = \frac{1}{N} \left| \sum_{i=1}^N \exp(\iota \theta_i) \right|, \quad (1)$$

where N is the number of individuals in the frame, ι is the imaginary unit, and θ_i is the heading of each individual.

Nearest-neighbour distance (NND) measures the level of “clustering” in a crowd. The average NND for a single “frame” (derived from either the real dataset or the simulation) is calculated from the sum of nearest-neighbour distances of all N individuals:

$$\nu = \frac{1}{N} \sum_{i=1}^N d_i, \quad (2)$$

381 where d_i is the nearest neighbor distance between point i and the closest individual in the
 382 frame, as calculated by the standard distance formula,

$$d_i = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}. \quad (3)$$

383 We selected 20 random Edinburgh clips with varying crowd sizes, and then simulated each
 384 scenario 20 times with each movement model. Results are presented in Figure 5; these
 385 confirm that both movement models produce high-level outputs that are comparable to
 386 the real-world scenarios, and that there are no significant differences between the outputs
 387 of each movement model.

388 **4.4 Classification tests**

389 For both tests, we constructed a web-based application¹ which presented users with an
 390 information screen, asked them to click to confirm their consent to participate, and then
 391 presented participants with a randomised sequence of movies. For each movie, partic-
 392 ipants were asked to click either a “Real” or “Simulated” button, according to their own
 393 perception and opinion. At the end of the sequence, users were asked in a free text box
 394 to supply short notes on any features that they thought allowed them to identify the real
 395 crowd, to specify their level of expertise in crowd science (“High”, “Medium” or “Low”), and
 396 to supply their email address (this was used as a participant ID to allow for tracking across
 397 the two tests). Once the user submitted their information, their responses were stored on
 398 the server, and they were told how many real crowds they had correctly identified (this may
 399 have inadvertently helped with recruitment, as some particularly high-scoring participants

¹Available at <http://www.martynamos.org/TTFC2/>

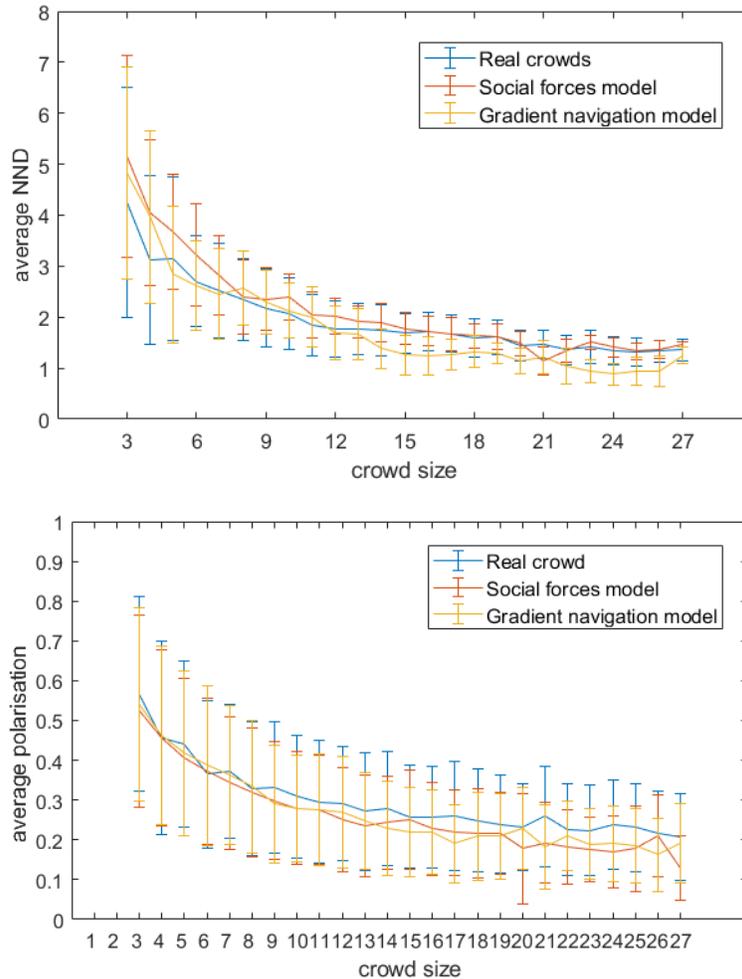


Figure 5: Movement models/real crowd statistical comparisons: Nearest Neighbour Distance (NND) (top) and polarisation (bottom) as a function of crowd size. The outputs of both movement models have properties that are close to those of the real crowds.

400 shared screenshots of their success on social media...)

401 **4.4.1 Test 1**

402 This was the “baseline” test to give each participant an initial score of their ability to classify
 403 movies as either real or simulated. We showed participants a sequence of 12 movies, 6 of
 404 which were based on real trajectories, and 6 of which were generated using the SFM-based
 405 simulation of that scenario. Each movie was 30s in duration (in all cases, participants were
 406 free to choose “early”, before the end of the movie, and move on to the next one).

407 For each real clip, the total number of individuals observed and average entry time interval
 408 is shown in Table 2 (the simulations were set up to reflect these). We present heatmap
 409 visualisations of the route choice distribution for each clip in Figure 6. The forum has 11
 410 ingress points, and the 12th row and column represent individuals who start or end their
 411 observed trajectories *inside* the forum space.

Clip	Number of individuals	Mean entry time interval (s)	Standard deviation (s)
1	194	0.34	0.22
2	149	0.46	0.26
3	112	0.67	0.38
4	104	0.62	0.34
5	150	0.48	0.24
6	125	0.55	0.33

Table 2: The total number of individuals observed and mean entry time interval of each clip from Test 1.

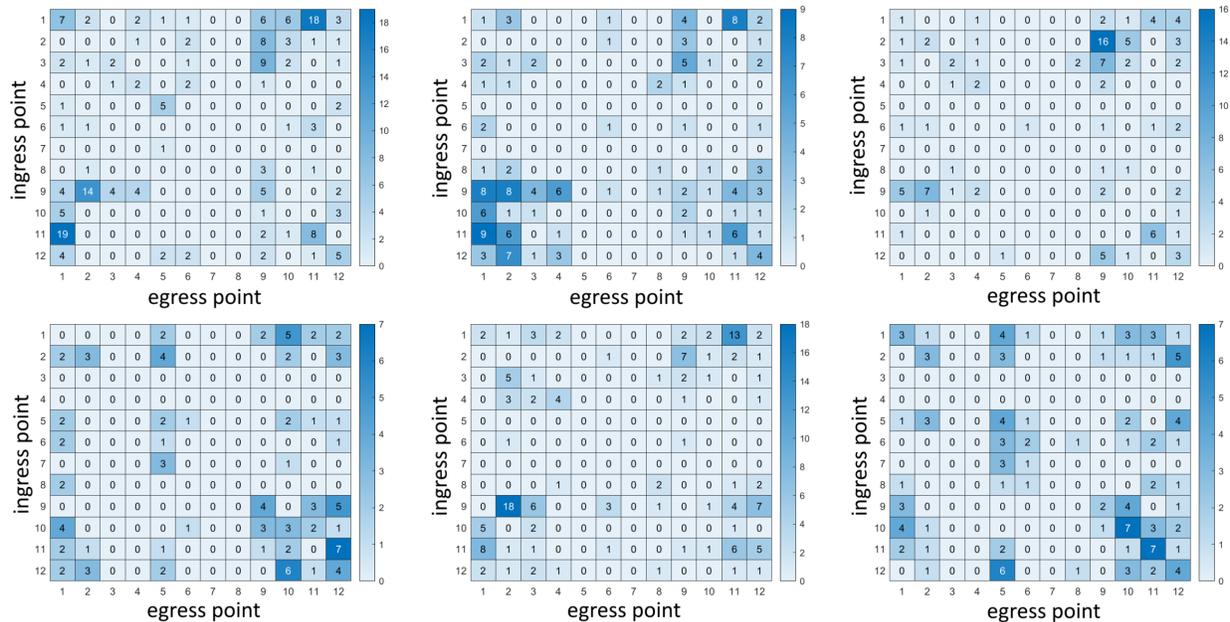


Figure 6: Heatmap representations of entry/exit point distributions for clips 1-3 (top) and 4-6 (bottom) from Test 1.

412 4.4.2 Test 2

413 We first required participants to undertake a training phase, in which they were shown
 414 6 representative clips generated from Edinburgh observations. Participants were made

Clip	Number of individuals	Mean entry time interval (s)	Standard deviation (s)
1	149	0.49	0.27
2	122	0.54	0.28
3	132	0.47	0.26
4	162	0.38	0.24
5	144	0.39	0.26
6	133	0.47	0.47

Table 3: The total number of individuals observed and mean entry time interval of each clip from Test 2.

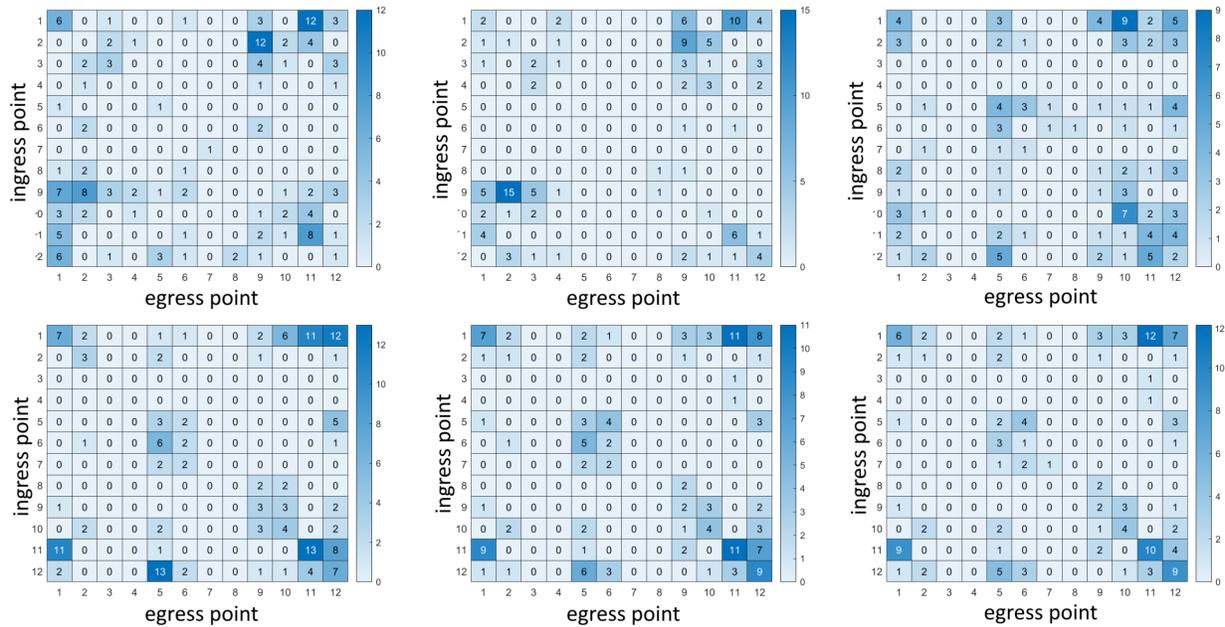


Figure 7: Heatmap representations of entry/exit point distributions for clips 1-3 (top) and 4-6 (bottom) from Test 2.

415 explicitly aware that they were watching real crowds. They were then shown 18 movies in
 416 total; 6 based on observations, 6 derived from SFM-based simulations, and 6 from GNM-
 417 based simulations.

418 For each real clip, the total number of individuals observed and average entry time interval
 419 is shown in Table 3 (again, the simulations were set up to reflect these). We present
 420 heatmap visualisations of the route choice distribution for each clip in Figure 7.

Set	Test 1	s.d.
$P_1 - P_2$	31.21%	20.19%
P_2	27%	19.31%

Table 4: Test 1 average scores for $P_1 - P_2$ and P_2 . Scores are presented as “% correctly classified”, as the number of movies differed between tests. Analysis confirms that P_2 is representative.

5 Results

In this Section we present our trial results. In what follows, we adopt the following notation for participant groups; P_1 is the initial set of 232 participants who took Test 1 (to establish their baseline scores, with no training) and P_2 is the subset of 50 participants in P_1 who went on to take Test 2 (the new test that included a training phase to establish whether or not performance improves after viewing real crowd videos).

5.1 Classification accuracy

We first consider whether or not group P_2 is representative of the larger set of participants. In both Test 1 and Test 2, participants were scored according to their ability to correctly classify movies, and received 1 point for every correct classification. We calculate the average Test 1 scores for both $P_1 - P_2$ (that is, participants who only took Test 1) and P_2 (participants who took both Tests), and present them in Table 4 (scores are presented as % due to the fact that the number of movies differed between tests).

A Lilliefors test confirms that neither dataset is normally distributed, so we use a two-sided Wilcoxon rank sum test to confirm that data in $P_1 - P_2$ and P_2 are samples from continuous distributions with equal medians ($p = 0.0724$). We conclude, therefore, that P_2 is a representative group.

We then calculate the average Test 1 and Test 2 classification scores for P_2 only; these are shown in Table 5. This reveals a *significant* improvement in overall correct classification score after training (from 27% to 60%). In Trial 2, participants correctly identified SFM-

Test 1	s.d.	Test 2	s.d
27%	19.31%	60.22%	26.35%

Table 5: Test 1 and Test 2 average scores for P_2 only.

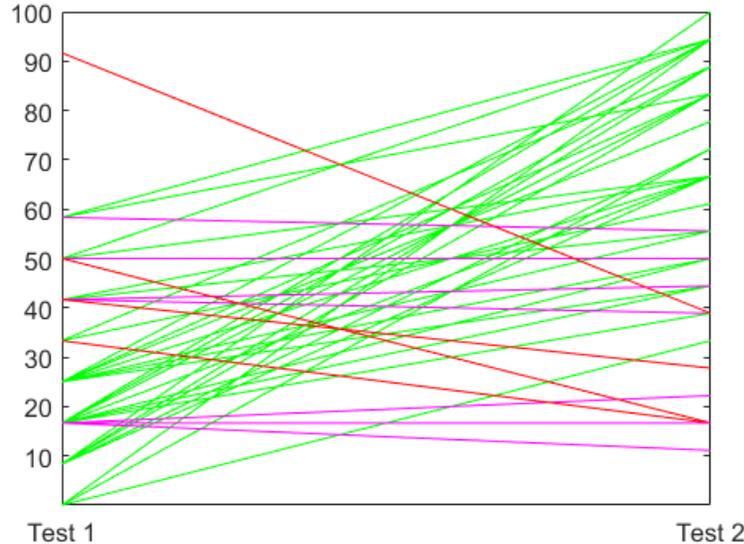


Figure 8: Slopegraph plot of changes in individual classification performance between Test 1 and Test 2 (50 individuals shown in total). Green lines show significant improvements, purple lines show small changes, and red lines show significant reductions in performance.

441 derived movies 63% of the time, and GNM-derived movies 59% of the time, so we cannot
 442 say that there exists a significant difference between the two models in terms of the overall
 443 characteristics of their outputs.

444 In Figure 8 we depict the individual changes in performance for the 50 members of P_2 ;
 445 visual inspection alone confirms that the vast majority of participants showed a marked
 446 improvement in classification performance after training. The average absolute change
 447 between Test 1 and Test 2 was 33.22%. If the participants had guessed at random in each
 448 test we would expect an average absolute change of 0%. A two-sided Wilcoxon signed
 449 rank test rejects the null hypothesis of a zero median in the distribution of average absolute
 450 change in our participant's test scores ($p < 0.001$). In Figure 9 we show the direction of
 451 improvement, confirming the bias towards an increase.

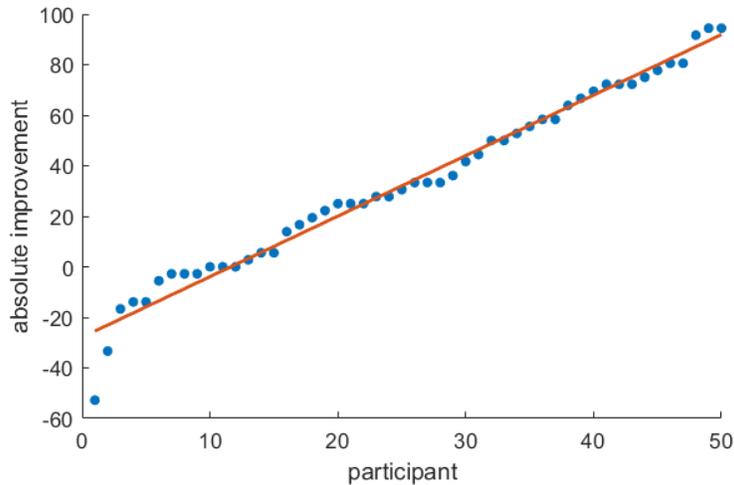


Figure 9: Trendline of absolute performance changes between Test 1 and Test 2 for P_2 participants.

452 These results confirm the first part of our hypothesis; that suitably trained individuals im-
 453 prove their classification performance after viewing movies of real crowds.

454 5.2 Narrative findings

455 We now move on to consider the free text supplied by members of P_2 , and extract com-
 456 mon themes that enable us to identify specific features of real and simulated crowds. We
 457 performed an initial version of this analysis in (Webster & Amos, 2020), but extracted only
 458 a small number of general themes, and did not correlate them with classification perfor-
 459 mance (as we do here). Our informal hypothesis is that participants who demonstrate
 460 significantly improved performance will correctly identify (in their free text responses) the
 461 characteristic features of both real and simulated crowds.

462 All 50 participants supplied feedback, so this provides useful additional context to explain
 463 the general uplift in performance. Given the relatively small amount of text, we performed
 464 manual thematic analysis to extract the predominant features highlighted in the supplied
 465 corpus. Each line of free text was broken down into thematic “atoms”, which were then
 466 semantically mapped onto over-arching themes. These are summarised in Table 6, par-

Real crowds	Freq. %
Heterogenous/diverse paths/speeds (R1)	9.21
Chaotic/unpredictable/erratic movement - rapid changes (R2)	21.05
Decisiveness/purposefulness - direct movement (R3)	6.56
Stop-start movement (R4)	7.89
Static individuals/groups (R5)	2.63
Groups/flocking/close proximity/collisions (R6)	7.89
Collision avoidance (R7)	5.26
Simulated crowds	Freq. %
Homogeneous behaviour (S1)	5.26
Rapid direction/speed changes (S2)	3.95
Goal-driven (S3)	3.95
Smooth/continuous movement (S4)	15.79
Clusters (S5)	1.32
Long interactions/collisions and close proximity (S6)	6.58
Collision avoidance (S7)	2.63

Table 6: Themes identified in narrative comments (labels given in brackets), and their observed frequencies. Related themes across “real” and “simulated” are numbered similarly, although there may not always be an exact correlation.

467 titioned into those features ascribed to real crowds, and those to simulated crowds. We
468 also give the relative frequency of each feature/theme (a link to the full dataset is supplied
469 at the end of the paper). We label each feature for ease of presentation/discussion.

470 We immediately notice two dominant features; R2 (*real* crowds exhibit chaotic or unpre-
471 dictable movement, sometimes with rapid changes in speed/direction) accounted for 21%
472 of thematic atoms, and S4 (*simulated* crowds show smooth/continuous movement) ac-
473 counted for nearly 16% of all atoms. These observations are clearly complementary, in
474 that (after training) observers believe that real crowds are more unpredictable than simu-
475 lated crowds, which move more smoothly. The real dataset does include many examples of
476 unpredictable/rapid changes in movement, where (we assume, not having access to the full
477 video datasets) an individual is “dashing” across the space and adjusting their movements
478 to avoid others, or where they double-back on themselves.

479 However, it is not sufficient to simply analyse the *frequency* of themes, since dominant
480 features may not necessarily correlate with good classification performance in the partici-

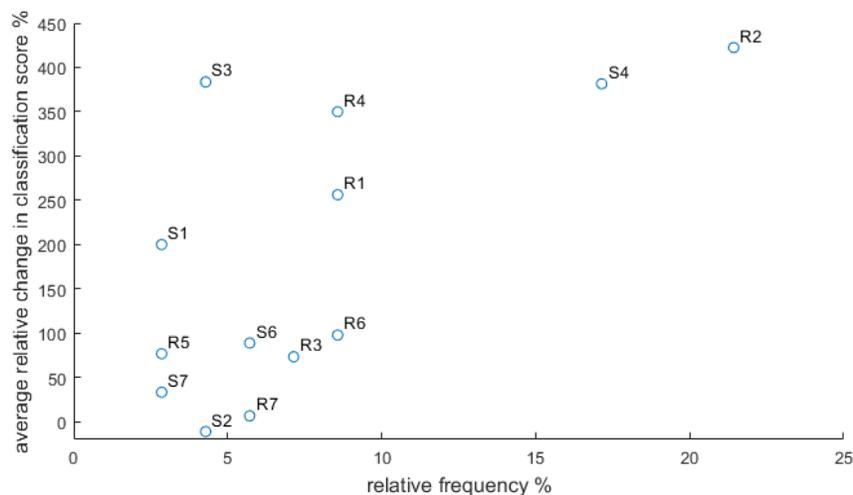


Figure 10: Thematic frequency versus average *relative* change in classification performance. The upper-right quadrant shows two themes (S4 and R2) which both appear frequently and which are correlated with significant positive relative change in classification performance in those participants who mention those themes.

481 pants who identify them. We also need to extract the features that have been identified by
 482 *the participants who perform best* (or who show the best relative improvement) in the clas-
 483 sification task. We first consider *relative* changes in scores, and then look at the *absolute*
 484 changes, as each perspective yields insights.

485 In Figure 10 we plot each theme against both their frequency of mentions and the average
 486 relative change in classification performance of participants who specifically mention that
 487 theme. All scores are expressed in terms of the *percentage* of movies that were correctly
 488 classified, not the “raw” score (as previously stated, the number of movies differs between
 489 tests). For each participant, only where $score_1 > 0$, the relative change in *score* is calculated
 490 by $((score_2 - score_1)/score_1 * 100)$. For example, a participant who scored 3/12 (25%) in
 491 Test 1 and 15/18 (83%) in Test 2 would have their relative change calculated as $((83 -$
 492 25)/25) * 100) = 232%.

493 When calculating the average relative change, we discard 4 participants with a Test 1 score
 494 of zero, as the notion of relative change is not defined for a zero reference value (however,
 495 these participants are still included in the discussion of actual score differences, below).

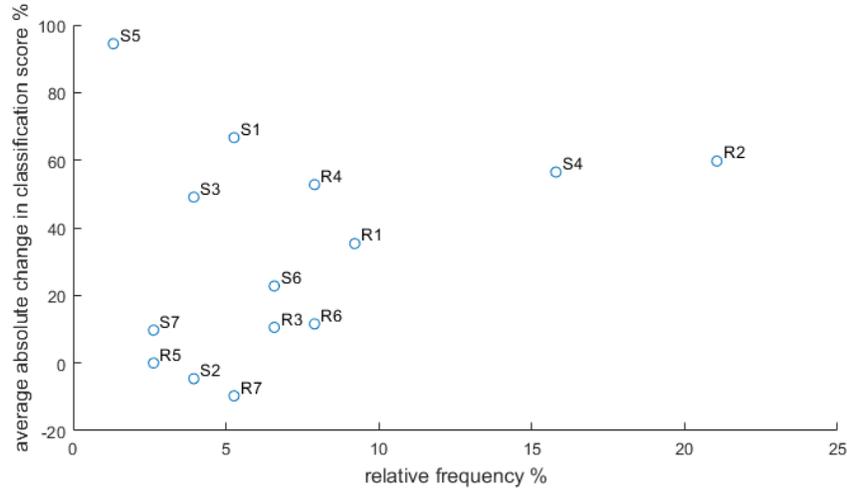


Figure 11: Thematic frequency versus average *absolute* change in classification performance. S2 and R7 are low -frequency themes that are nonetheless associated with reductions in classification performance.

496 We notice, from inspection, a cluster of themes that are relatively infrequently mentioned
 497 (< 10%), but which are associated with significant improvements in classification perfor-
 498 mance. However, we see that the two themes that are mentioned with frequency > 15% -
 499 S4 (smooth/continuous movement in simulated crowds) and R2 (unpredictable movement
 500 in real crowds) - are both also associated with performance improvements of around 400%.
 501 As noted earlier, these themes are complementary.

502 This finding is entirely consistent with our earlier informal narrative results (Webster &
 503 Amos, 2020), where participants who had “flipped” the real and simulated crowds believed
 504 that erratic movement was characteristic of “fake” (simulated) crowds, and that real crowds
 505 moved smoothly and predictably. After training on real crowds, however, the participants
 506 in this second trial correctly identified that real crowds are actually more noisy and unpre-
 507 dictable, and that overwhelmingly smooth, predictable trajectories are a characteristic of
 508 simulations.

509 We now consider *absolute* changes in classification score between tests. We see roughly
 510 the same clustering of labels as before (S5: presence of clusters in simulated crowds is

511 an outlier, in that it was mentioned only by a single person, albeit one who saw a sig-
512 nificant improvement in their classification score). Here we draw particular attention to
513 the (albeit infrequently mentioned) themes that are correlated with *negative* shifts in per-
514 formance. That is, the features that are mentioned by participants whose classification
515 performance got worse after training. The two features to which this applies are S2 (rapid
516 direction/speed changes in simulated crowds) and R7 (collision avoidance in real crowds).
517 Again, these findings are entirely consistent with both the current results and our previous
518 study. If high-performing participants correctly spot that simulated crowds move smoothly,
519 then it is entirely to be expected that low-performing participants will (incorrectly) ascribe
520 S2 to them. Collision avoidance in real crowds (R7) is also specifically mentioned in our
521 previous study; participants who performed badly assumed that individuals in real crowds
522 would naturally avoid one another. As we observe in (Webster & Amos, 2020), “In reality,
523 the opposite is true, as the real dataset contains multiple instances of individuals coming
524 into close proximity. Moreover, the social forces model explicitly tries to keep individu-
525 als apart unless close proximity is unavoidable, so the behaviour (distance keeping) that
526 participants attributed to real people was actually an in-built feature of the simulation.”
527 However, we must approach these findings with a degree of caution, as it may be the
528 case (for example) that the high-performing individuals are simply better learners, or some
529 videos may be inherently easier (or more difficult) to classify. All we claim here is that
530 there would appear to be a *correlation* between high classification performance and a
531 small set of identifiable features of crowds. An investigation of the fundamental underlying
532 process(es) is beyond the scope of the current paper, but may be performed in future work.
533 Based on these findings, we conclude that the primary feature of real crowds that allows
534 trained individuals to correctly distinguish them from simulated crowds is their higher de-
535 gree of unpredictability in terms of individual trajectories. A secondary feature is collision
536 avoidance (specifically, proximity). Based on this work, our main suggestion (if what we

537 seek is realistic believability in crowd simulations) is that models should include the facility
538 to add a degree of unpredictability to the movement of individual agents (surprisingly, this
539 feature is not generally provided). Models might also benefit from a relaxation of collision
540 detection radii to allow for closer proximity of agents. In this way, we might easily replicate
541 the appearance of at least some of the micro-level behaviours referenced by (Lerner et al.,
542 2007).

543 **6 Discussion and Conclusions**

544 In this paper we report the results of a human trial to identify the “signature” characteristics
545 of real crowds that allow them to be distinguished from simulated crowds. We find that
546 unpredictability in terms of individual trajectories is by far the best discriminator, and
547 proximity in collision detection is also relevant. We note some limitations of our study;
548 the underlying crowd dataset is based on a relatively small physical space which is quite
549 regular in nature, but we point out that it is actually much larger than the arenas used for
550 artificial crowd experiments. Moreover, the observations have a higher level of ecological
551 validity, as the recorded pedestrians were not consciously aware of being participants in
552 an experiment. Our second test used a relatively small number of participants, but we
553 have established that they were representative of a larger set. Finally, our findings are
554 only applicable to “routine” crowds (that is, where people are going about their everyday
555 business), and not to “emergency” or “evacuation” crowds, where behaviours will be very
556 different.

557 However, there is still significant value in updating simulation of such routine crowds to
558 render them more realistically, especially if important policy or design decisions are to be
559 made based on how they are perceived. With this in mind, there may be value in training
560 decision-makers who use such simulations as part of their process (in a manner similar to
561 that performed in our Test 2), in order to ensure that they can first detect the characteristic

562 features of real crowds (as opposed to making decisions based on flawed assumptions of
563 how crowds behave). Fundamentally, the value of additional realism in crowd simulations
564 may only be realised if end-users are able to *recognise* it.

565 This study has provided empirical evidence to support the inclusion of relatively straight-
566 forward modifications to any and all of the movement models underpinning both scientific
567 and commercial crowd simulation packages. Importantly, the addition of noise to individ-
568 ual trajectories and the relaxation of collision detection radii are entirely generic updates,
569 but ones that could significantly improve the believability of crowd simulations across a
570 range of applications.

571 Future work may include the automatic detection of features of real crowds from larger
572 and more complex datasets, consideration of the impact of changing movement model
573 parameters, and the integration of identified features into commercial crowd simulation
574 packages in order to test their impact on believability (thus “closing the circle”).

575 **7 Materials**

576 All code (simulations and analysis scripts) and datasets generated are available at [http:](http://doi.org/10.6084/m9.figshare.c.5280902)
577 [//doi.org/10.6084/m9.figshare.c.5280902](http://doi.org/10.6084/m9.figshare.c.5280902)

578 **Acknowledgements**

579 JW was supported by a Ph.D. studentship from the Faculty of Engineering and Environ-
580 ment, Northumbria University. We thank Gerta Köster and her research team for useful
581 discussions, and all of the trial participants for their contributions.

References

- 582
- 583 Adrian, J., Amos, M., Baratchi, M., et al. (2019). A glossary for research on human crowd
584 dynamics. *Collective Dynamics*, 4(A19), 1–13.
- 585 Aschwanden, G., Haegler, S., Bosché, F., Van Gool, L., & Schmitt, G. (2011). Empiric design
586 evaluation in urban planning. *Automation in Construction*, 20(3), 299–310.
- 587 Bastien, R., & Romanczuk, P. (2020). A model of collective behavior based purely on vision.
588 *Science Advances*, 6(6), eaay0792.
- 589 Bode, N. W., Franks, D., & Wood, J. (2010). Making noise: Emergent stochasticity in collective
590 motion. *Journal of Theoretical Biology*, 267(3), 292–299.
- 591 Bode, N., Kemloh Wagoum, A., & Codling, E. (2015). Information use by humans during
592 dynamic route choice in virtual crowd evacuations. *Royal Society Open Science*,
593 2(1).
- 594 Bohannon, R. (1997). Comfortable and maximum walking speed of adults aged 20–79
595 years: Reference values and determinants. *Age and Ageing*, 26(1), 15–19.
- 596 Brambilla, M., Ferrante, E., Birattari, M., & Dorigo, M. (2013). Swarm robotics: A review from
597 the swarm engineering perspective. *Swarm Intelligence*, 7(1), 1–41.
- 598 Crociani, L., Lämmel, G., & Vizzari, G. (2016). Multi-scale simulation for crowd management:
599 A case study in an urban scenario. *International Conference on Autonomous Agents
600 and Multiagent Systems (AAMAS)*, 147–162.
- 601 Cronin, L., Krasnogor, N., Davis, B., Alexander, C., Robertson, N., Steinke, J., Schroeder, S.,
602 Khlobystov, A., Cooper, G., Gardner, P., et al. (2006). The imitation game - a compu-
603 tational chemical approach to recognizing life. *Nature Biotechnology*, 24(10), 1203.
- 604 Duives, D., Daamen, W., & Hoogendoorn, S. (2013). State-of-the-art crowd motion simula-
605 tion models. *Transportation Research Part C: Emerging Technologies*, 37, 193–209.
- 606 Dutra, T., Marques, R., Cavalcante-Neto, J., Vidal, C., & Pettré, J. (2017). Gradient-based
607 steering for vision-based crowd simulation algorithms. *Computer Graphics Forum*,
608 36, 337–348.

- 609 Ennis, C., Peters, C., & O'Sullivan, C. (2011). Perceptual effects of scene context and view-
610 point for virtual pedestrian crowds. *ACM Transactions on Applied Perception (TAP)*,
611 8(2), 1–22.
- 612 Feng, T., Yu, L.-F., Yeung, S.-K., Yin, K., & Zhou, K. (2016). Crowd-driven mid-scale layout
613 design. *ACM Transactions on Graphics*, 35(4), 132–1.
- 614 Fernando, T., Denman, S., Sridharan, S., & Fookes, C. (2018). Tracking by prediction: A deep
615 generative model for multi-person localisation and tracking. *2018 IEEE Winter Con-*
616 *ference on Applications of Computer Vision (WACV)*, 1122–1132.
- 617 Fuchsberger, A., Tahmasbi, N., & Ricks, B. (2017). A framework for achieving realism in
618 agent-based pedestrian crowd simulations. *AMCIS 2017 - America's Conference*
619 *on Information Systems*.
- 620 Gloor, C. (2016). PedSim: Pedestrian crowd simulation [<http://pedsim.silmaril.org>].
- 621 Gorochowski, T. (2016). Agent-based modelling in synthetic biology. *Essays in Biochem-*
622 *istry*, 60(4), 325–336.
- 623 Harding, P., Gwynne, S., & Amos, M. (2011). Mutual information for the detection of crush.
624 *PloS ONE*, 6(12), e28747.
- 625 Harel, D. (2005). A Turing-like test for biological modeling. *Nature Biotechnology*, 23(4),
626 495.
- 627 Helbing, D., & Molnar, P. (1995). Social force model for pedestrian dynamics. *Physical Re-*
628 *view E*, 51(5), 4282.
- 629 Herbert-Read, J., Romenskyy, M., & Sumpter, D. (2015). A Turing test for collective motion.
630 *Biology Letters*, 11, 20150674.
- 631 Kimura, T., Sekine, H., Sano, T., Takeichi, N., Yoshida, Y., & Watanabe, H. (2003). Pedestrian
632 simulation system SimWalk. *Summaries of Technical Papers of the Annual Meeting*
633 *of the Architectural Institute of Japan, E-1*, 915–916.

- 634 Klüpfel, H. (2007). The simulation of crowd dynamics at very large events - calibration,
635 empirical data, and validation. *Pedestrian and Evacuation Dynamics (PED) 2005*
636 (pp. 285–296). Springer.
- 637 Korhonen, T., Hostikka, S., Heliövaara, S., & Ehtamo, H. (2010). FDS+Evac: An agent based
638 fire evacuation model. *Pedestrian and Evacuation Dynamics (PED) 2008* (pp. 109–
639 120). Springer.
- 640 Lemercier, S., & Auberlet, J. (2016). Towards more behaviors in crowd simulation. *Computer*
641 *Animation And Virtual Worlds*, 27(1), 24–34.
- 642 Lerner, A., Chrysanthou, Y., & Lischinski, D. (2007). Crowds by example. *Computer Graphics*
643 *Forum*, 26(3), 655–664.
- 644 Loscos, C., Marchal, D., & Meyer, A. (2003). Intuitive crowd behavior in dense urban environ-
645 ments using local laws. *Proceedings - Theory and Practice of Computer Graphics,*
646 *TPCG 2003*, 122–129.
- 647 Lovreglio, R., Dias, C., Song, X., & Ballerini, L. (2017). Towards microscopic calibration of
648 pedestrian simulation models using open trajectory datasets: The case study of the
649 Edinburgh Informatics Forum. *Conference on Traffic and Granular Flow, Washington*
650 *DC, USA*.
- 651 Luo, L., Zhou, S., Cai, W., Yoke, M., Low, H., Tian, F., Wang, Y., Xiao, X., & Chen, D. (2008).
652 Agent-based human behavior modeling for crowd simulation. *Computer Animation*
653 *And Virtual Worlds*, 19(August), 271–281.
- 654 Mahmood, I., Haris, M., & Sarjoughian, H. (2017). Analyzing emergency evacuation strate-
655 gies for mass gatherings using crowd simulation and analysis framework: Hajj sce-
656 nario. *Proceedings of the 2017 ACM SIGSIM Conference on Principles of Advanced*
657 *Discrete Simulation*, 231–240.
- 658 Majecka, B. (2009). *Statistical models of pedestrian behaviour in the forum* (Master's the-
659 sis). School of Informatics, University of Edinburgh.

660 Mckenzie, F., Petty, M., Kruszewski, P., Gaskins, R., Nguyen, Q.-A., Seevinck, J., & Weisel, E.
661 (2008). Integrating crowd-behavior modeling into military simulation using game
662 technology. *Simulation & Gaming*, 39(1), 10–38.

663 Moussaïd, M., Kapadia, M., Thrash, T., Sumner, R., Gross, M., Helbing, D., & Hölscher, C.
664 (2016). Crowd behaviour during high-stress evacuations in an immersive virtual en-
665 vironment. *Journal of The Royal Society Interface*, 13(122), 20160414.

666 Pelechano, N., Allbeck, J. M., & Badler, N. I. (2007). Controlling individual agents in high-
667 density crowd simulation. *Proceedings of the 2007 ACM SIGGRAPH Eurographics*
668 *symposium on Computer animation*, 1, 108.

669 Peters, C., & Ennis, C. (2009). Modeling groups of plausible virtual pedestrians. *IEEE Com-*
670 *puter Graphics and Applications*, 29(4), 54–63.

671 Pettré, J., Ondřej, J., Olivier, A.-H., Cretual, A., & Donikian, S. (2009). Experiment-based
672 modeling, simulation and validation of interactions between virtual walkers. *Pro-*
673 *ceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer An-*
674 *imation*, 189–198.

675 Pouw, C., Toschi, F., van Schadewijk, F., & Corbetta, A. (2020). Monitoring physical dis-
676 tancing for crowd management: Real-time trajectory and group analysis. *PloS ONE*,
677 15(10), e0240963.

678 Pretorius, M., Gwynne, S., & Galea, E. (2015). Large crowd modelling: An analysis of the
679 Duisburg Love Parade disaster. *Fire and Materials*, 39(4), 301–322.

680 Reynolds, C. (1987). Flocks, herds and schools: A distributed behavioral model. *Proceedings*
681 *of the 14th Annual Conference on Computer Graphics and Interactive Techniques*,
682 25–34.

683 Ricks, B. (2013). *Improving Crowd Simulation with Optimal Acceleration Angles, Move-*
684 *ment on 3D Surfaces, and Social Dynamics* (Doctoral dissertation). Brigham Young
685 University.

- 686 Rivers, E., Jaynes, C., Kimball, A., & Morrow, E. (2014). Using case study data to validate 3D
687 agent-based pedestrian simulation tool for building egress modeling. *Transporta-*
688 *tion Research Procedia*, 2, 123–131.
- 689 Seer, S., Rudloff, C., Matyus, T., & Brändle, N. (2014). Validating social force based models
690 with comprehensive real world motion data. *Transportation Research Procedia*, 2,
691 724–732.
- 692 Seitz, M., Templeton, A., Drury, J., Köster, G., & Philippides, A. (2017). Parsimony versus
693 reductionism: How can crowd psychology be introduced into computer simulation?
694 *Review of General Psychology*, 21(1), 95–102.
- 695 Singh, H., Arter, R., Dodd, L., Langston, P., Lester, E., & Drury, J. (2009). Modelling sub-
696 group behaviour in crowd dynamics DEM simulation. *Applied Mathematical Mod-*
697 *elling*, 33(12), 4408–4423.
- 698 Smith, A., James, C., Jones, R., Langston, P., Lester, E., & Drury, J. (2009). Modelling contra-
699 flow in crowd dynamics DEM simulation. *Safety Science*, 47(3), 395–404.
- 700 Templeton, A., Drury, J., & Philippides, A. (2015). From mindless masses to small groups :
701 conceptualizing collective behavior in crowd modeling. *Review of General Psychol-*
702 *ogy*, 19(3), 215–229.
- 703 Thalmann, S., & Musse, S. (2013). *Crowd Simulation*. Springer.
- 704 Turing, A. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433.
- 705 Wagner, N., & Agrawal, V. (2014). An agent-based simulation system for concert venue
706 crowd evacuation modeling in the presence of a fire disaster. *Expert Systems with*
707 *Applications*, 41(6), 2807–2815.
- 708 Wang, H., & O’Sullivan, C. (2016). Globally continuous and non-Markovian crowd activity
709 analysis from videos. *European Conference on Computer Vision*, 527–544.
- 710 Wang, W. L., Lo, S. M., & Liu, S. B. (2017). A cognitive pedestrian behavior model for ex-
711 ploratory navigation: Visibility graph based heuristics approach. *Simulation Mod-*
712 *elling Practice and Theory*, 77, 350–366.

- 713 Webster, J., & Amos, M. (2020). A Turing test for crowds. *Royal Society Open Science*,
714 7(200307).
- 715 Wei, X., Lu, W., Zhu, L., & Xing, W. (2018). Learning motion rules from real data: Neural
716 network for crowd simulation. *Neurocomputing*, 310, 125–134.
- 717 Yao, Z., Zhang, G., Lu, D., & Liu, H. (2020). Learning crowd behavior from real data: A residual
718 network method for crowd simulation. *Neurocomputing*, 404, 173–185.
- 719 Zhang, D., Zhu, H., Hostikka, S., & Qiu, S. (2019). Pedestrian dynamics in a heterogeneous
720 bidirectional flow: Overtaking behaviour and lane formation. *Physica A: Statistical
721 Mechanics and its Applications*, 525, 72–84.
- 722 Zönnchen, B., Kleinmeier, B., & Köster, G. (2020). Vadere – a simulation framework to com-
723 pare locomotion models. *Traffic and Granular Flow 2019* (pp. 331–337). Springer.