Representative Synthetic Crowds for Inclusive Environment Design

Brandon Haworth, Mubbasir Kapadia, Petros Faloutsos

2021

The final publication is available at:

https://doi.org/10.1109/AIVR52153.2021.00035

Citation for this paper:

Haworth, B., Kapadia, M., & Faloutsos, P. (2021). Representative Synthetic Crowds for Inclusive Environment Design. 2021 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR), 150–153. https://doi.org/10.1109/AIVR52153.2021.00035

# Representative Synthetic Crowds for Inclusive Environment Design

*Abstract*—**Synthetic crowds serve as a powerful tool for numerous applications across industry and research. We argue that, while synthetic crowds may also be a highly valuable tool for predictive and inclusive building design, the current state-of-the-art is not representative enough to safely impact this space. We show that model fidelity and non-biomechanical rules can not capture the non-linear kinematics of both normative and non-normative gaits. Finally, we argue that recent developments in deep reinforcement learning may afford a significant increase in fidelity and a move away from limited data driven methods, ad hoc or expert rules, and heuristics. These new approaches, however, also have several issues that are the focus of current research and could one day serve as the groundwork for high fidelity inclusive design processes driven by simulation.**

*Index Terms*—**Synthetic Crowds; Human Movement Simulation; Inclusive Environment Design**

## I. Introduction

The modern world is increasingly crowded. As global population rises, and factors like the portion of the ageing population increases, it is important to consider more predictive and inclusive approaches to environment design. Our cities are made up of necessarily diverse peoples. One of the dimensions of diversity directly impacted by the design of our environments is mobility. Several countries have made strong moves towards codifying accessibility in the built environment. For example, in the United States, the Americans with Disabilities Act of 1990 (ADA) codifies several vectors of accessibility requirements, forming a prescriptive system. Similarly, in Canada, provincial-level acts like the Accessibility for Ontarians with Disabilities Act (AODA), or at the federal level, the not yet realized, Accessible Canada Act (ACA), aim to create and enforce accessibility standards. However, despite somewhat united efforts, including the 2010 United Nations Convention on the Rights of Persons with Disabilities, accessibility in practice has been problematic. Accessibility is challenged by very slow changes, attrition in adoption, disconnects in information gathering and availability, and lack of enforcement [1]. These codes and approaches are either too prescriptive, or too performance-based, struggle to capture edge cases, and are typically defined in static cases for limited definitions of disability. The environment design pipeline that results from this often seeks to meet minimum requirements – design comes first and code audits must simply be passed. Synthetic crowds have, over the last decade, emerged as a means of dynamically analysing building, environment, and event designs. Having roots in graphics and animation, synthetic crowds offer a means of large scale testing through simulation that is both cost and

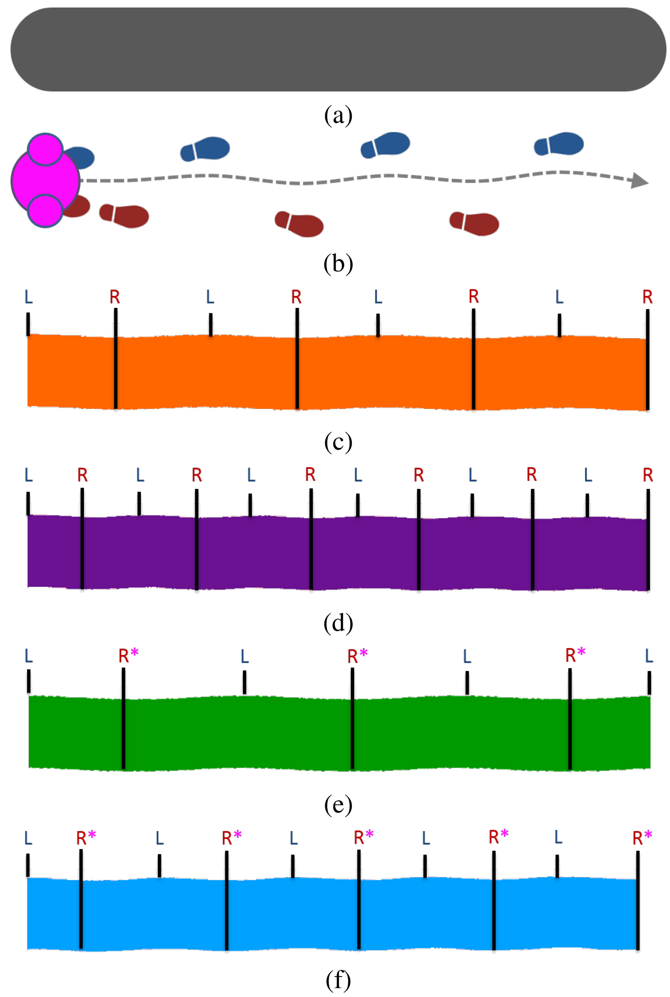


Fig. 1. Collision corridors of multiple models with left (L), right (R) foot flat, and paretic side (*) conditions annotated for post stroke models. Differences in collision corridor shape are apparent between (a) the single-particle model's collision corridor and (b) the footstep model's (c) normative walking collision corridor, as well as the collision corridors for (d) low-velocity walking, (e) temporal gait asymmetry post-stroke and (e) low-velocity temporal gait asymmetry post-stroke. The single-particle model can not capture the non-linear corridor produced by the biomechanics in the footstep model, and the low-velocity proxy can not capture the asymmetry in either of the post-stroke models.

time effective. More recently, automated methods driven by optimization and machine learning have emerged which afford rapid exploration of many environments and scenarios using synthetic crowds. That is, synthetic crowds afford predictive analysis of environments, and that analysis can feed predictive decision-making. However, our work shows that synthetic crowds are not representative. We show that current state-of-the-art representativeness relies on simplified representation of individuals and underlying heuristics or rules determine the movement.

This paper posits that if models are not representative then designs following the use of synthetic crowds are not inclusive. Our work seeks to mitigate the representativeness of synthetic crowds with more sophisticated body modelling and model-free parametric policy learning. In this paper, we frame the work with the built environment's role in the disablement process. We then explore the impact of model representation and diverse biomechanics on synthetic crowds. Finally, discuss model-free data-driven machine learning, issues related to this approach, and how it may hold the answer to learning more representative models. Several recent developments in deep reinforcement learning, including parametric policy learning, combinatorial curriculum learning, and hierarchical learning, hold promise for developing new approaches with high-fidelity representation and generalizable policies for diverse agent-based simulation.

## II. Built Environment and Disability

There are several notable models of disabilities, and all can be framed within the domain of the built environment [1]. Relevant to this work are two particular models, the medical model and the social model. The medical model pathologizes illness, disability, and difference, while the social model centres on the humanity and autonomy of the person or people of focus. This places the burden of concepts such as illness, disability, and difference on environment and the organization of society. In this way, these concepts are developed and reinforced by the institutionalization and regulation of bodies [2]. The social model explicitly places the onus of the disablement process on society, its organizations, and those assumptions are reflected in the built environment [3]. The impacts of the various components in this problem space are of interest to the many areas of critical disability studies, in particular the role of environment design in the disablement process and its impact on mobility, health, and socio-economic factors of community members with disabilities [4]–[6]. Many elements of the disablement process specifically intersect with mobility, or lack thereof, within community and city environments.

## III. Importance of Model and Representation

This section addresses the notion of fidelity and the use of proxies for diversity in modelling. The following two experiments separate the level at which biomechanical heterogeneity impacts steering in synthetic crowds: the collider corridor formed by moving an agent's colliders through space-time; and the crowd outcome measures used to understand scenarios.

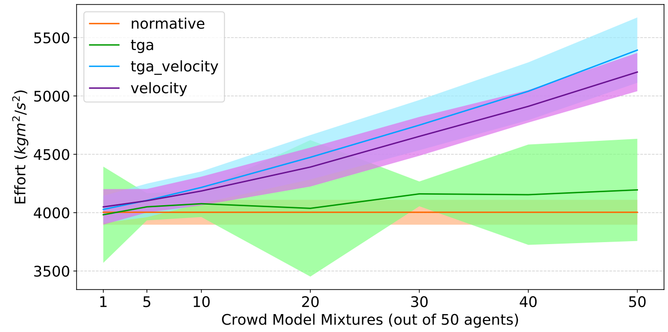


Fig. 2. Fundamental diagram of total crowd effort over model mixtures. The lines and shaded areas represent the means ($\mu$) and standard deviations ($\sigma$) respectively. Results show significant differences in overall crowd effort between normative walking, velocity-only proxies for diversity (velocity), diverse temporal gait asymmetry (tga), and diverse temporal gait asymmetry with velocity control (tga_velocity) in the context of the common bottleneck egress scenario (emergency exit). This serves as evidence of velocity-only proxies for locomotion heterogeneity (such as diversity in gait) being insufficient for representing diverse populations in synthetic crowds.

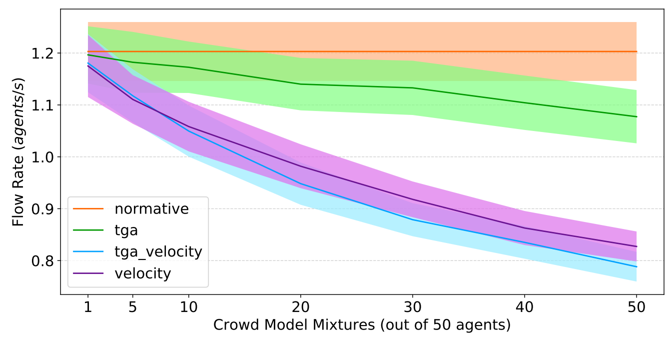


Fig. 3. Fundamental diagram of total crowd flow rate over model mixtures. The lines and shaded areas represent the means ($\mu$) and standard deviations ($\sigma$) respectively. Results show a significant differences in overall crowd flow rate between normative walking, velocity-only proxies for diversity (velocity), diverse temporal gait asymmetry (tga), and diverse temporal gait asymmetry with velocity control (tga_velocity) in the context of the common bottleneck egress scenario (emergency exit). This serves as evidence of velocity-only proxies for locomotion heterogeneity (such as diversity in gait) being insufficient for representing diverse populations in synthetic crowds.

Post-stroke temporal gait asymmetry (TGA), is used as a probe for exploring non-linear biomechanical heterogeneity. We also examine velocity as a common proxy for disordered gaits by modelling both a low-velocity condition (similar to Parkinsonian shuffling), and a TGA low-velocity condition (observed in the literature). The gait conditions and normative walking are implemented using a high fidelity multi-particle footstep-based model for synthetic crowds [7], [8].

*1) Impact of Fidelity and Diverse Biomechanics on Collisions:* The corridors of each model in a straight walking task are recorded for all gait conditions and a single-particle model. The single-particle model is a baseline for typical low fidelity models that represent the entire human as a circle in a point-mass system. The implication here is that if an agent's collider corridor shape is predicated on its underlying model and representation, and that shape dictates local interactions with other agents, then changes in shape would alter interacting agents steering behaviours. The qualitative difference between the models' collider corridors is highlighted in Figure 1 by demarcating the approximate left and right foot flat, or loading response phase, conditions. The normative model produces a regular symmetric stepping pattern where the mean stride length is $0.762$m, step lengths are then approximately $0.381$m, and spatial stride frequency is approximately $1.32$Hz. The velocity model produces a regular symmetric pattern of stride lengths approximately $0.508$m, step lengths of approximately $0.254$m), and spatial stride frequency is approximately $2.068$Hz. The TGA model produces asymmetrical stepping pattern with shorter paretic side step. The mean length of paretic side steps is $0.431$m, while the for non-paretic side the mean length is $0.504$m. The spatial stride frequency is approximately $0.9398$Hz. The TGA velocity model produces asymmetrical stepping pattern with shorter paretic side step. The mean length of paretic side steps is $0.28$m, while the for non-paretic side the mean length is $0.299$m. The spatial stride frequency is approximately $1.6917$Hz.

Both qualitative analysis and step length, stride length, and spatial stride frequency measures reveal that none of the traversed collidor corridors is like the other. It is clear in this study that the shape of the collider corridors is impacted both by the underlying agent representation and by the modelling of biomechanical parameters.

*2) Impact of Diverse Biomechanics on Crowd Outcome Measures:* This study examines the impact of diversity in locomotion biomechanics on common crowd outcomes measures common in the analysis of environments. The hypothesis of this study, given the results of the collider corridor study, is that diversity in locomotion biomechanics within a crowd changes common crowd outcome measures at different levels of diversity. Here we use the commonly seen bottleneck egress scenario (one way hallway exit, or emergency exit) with $50$ total agents tasked with exiting. In this study, heterogeneity is introduced in a controlled manner. That is, the study samples multiple levels of heterogeneity, namely: $1/50$ ($2\%$), $5/50$ ($10\%$), $10/50$ ($20\%$), $20/50$ ($40\%$), $30/50$ ($60\%$), $40/50$ ($80\%$), $50/50$ ($100\%$); where the ratio is *agents with a non-normative model* divided by *total agents* with the remaining agents from the total using the normative effort optimal parametrization of the model [9].

The results of this study show significant issues with representational proxies in synthetic crowd outcomes. We check for significant differences with a Kruskal Wallis test and a battery of three post-hoc tests (Conover, Dunn, & Nemenyi) where all pairwise post-hoc tests are considered significant only if $p < 0.01$ across all three tests. Both effort, as seen in Figure 2, and flow rate, as seen in Figure 3, measures rapidly and significantly diverge across conditions at levels which would be seen in everyday environments and care facilities. With respect to effort expenditure: after $10/50$ ($20\%$), normative and TGA models; after $20/50$ ($40\%$), TGA low velocity and low velocity normative models; and other models at all levels differentiate significantly. As well TGA and even further, TGA low velocity, is more costly with respect to effort–an effect reflected in the literature [10]–[12]. With respect to flow rate: after $5/50$ ($10\%$), normative and both TGA and TGA low velocity; after $20/50$ ($40\%$), TGA low velocity and low velocity normative models; and other models at all levels differentiate significantly.

This is clear evidence that diversity in the representation of synthetic crowds agents matters and that velocity is not a proxy for modelling disability (here we extrapolate from the impact of temporal gait asymmetry that other disabilities or diverse mobilities are likely not well represented by simple proxies).

## IV. Representative Crowds

Past methods for synthetics crowds can broadly be placed into two groups: data-driven and rule-based. Data-driven methods tend to localize to the data. Rule-based methods can not account for a broad diversity of individuals. Advances in data-driven deep reinforcement learning (DRL) hold promise for learning diverse agent policies with much more adequate representations. However, there are three particular problems with DRL: data; heterogeneity/diversity; and generality. A key issue with data in any machine learning problem is representation. If the data is not representative or is poorly conceived then the learning will replicate these representation problems. With respect to heterogeneity/diversity, the policies underlying many simulators are largely homogeneous to make the crowd problem tractable. Finally, it is very difficult to build or learn a model that both performs well and is representative outside the scope of its construction or learning.

With respect to data, it is important to build new datasets where the focus is on gathering from largely under-represented groups in a way that is clean and reusable for others. This can be achieved by working specifically with people with gait disorders or mobility devices. Additionally, it is important to properly collect a broad set of gait specific outcome measures with grounding in the relevant research. This meta-information is often useful in rehabilitation and is likely useful in learning.

With respect to diversity, RL is largely based on history during learning, having multiple interacting agents who are progressively learning means that history is less usable/stable. This is known as the non-stationary problem–as an agent learns, other agents' history of how that agent acted is now less useful (the environment shifts). To address this, it is possible to construct a hierarchical learning problem such that tasks or skills can be learned constructively. The theory behind this hierarchical approach is that it improves exploration of the policy space [13], but it also provides the affordance that portions of the model, such as the lower-level control of locomotion, may be shared amongst agents [14], [15]. However, this means that the locomotion becomes fundamentally homogeneous, an avenue of current research–as seen in Figure 4. Additionally, it is beneficial to have diverse parametric agents that can be configured for a variety of purposes. To address this, it is possible to learn a large family of policies which are conditioned on such a parameter [16].

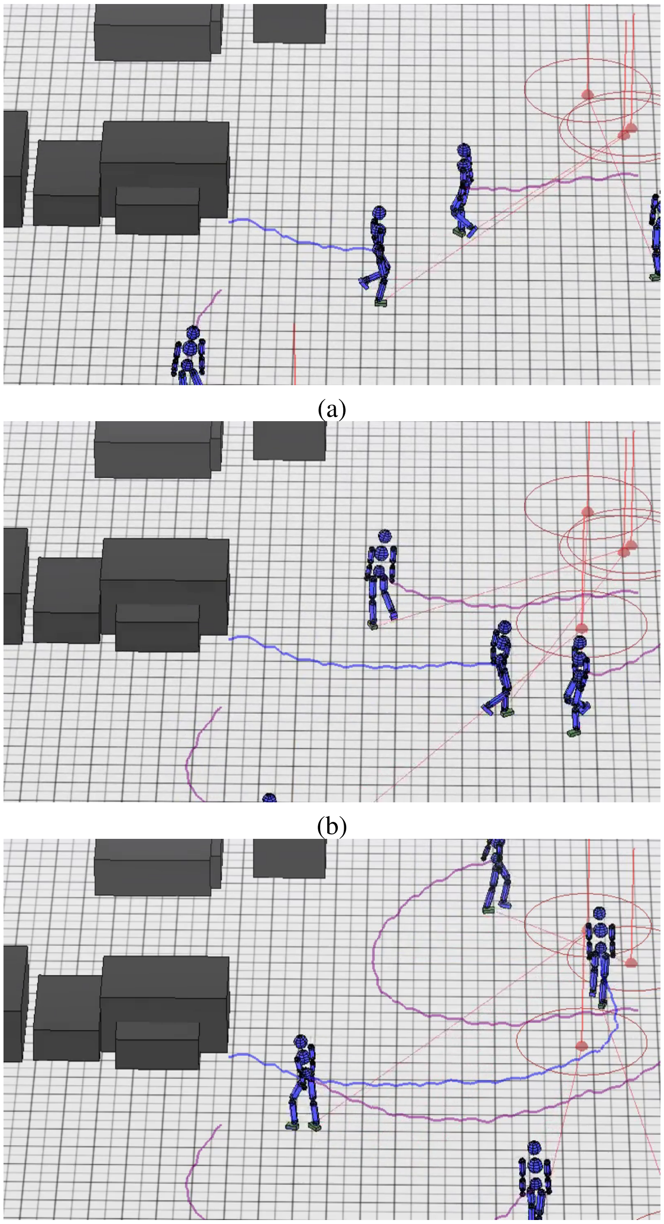

(a)

(b)

(c)

Fig. 4. Example of (a-c) rasterized frames of a video of full body articulated characters trained using a multi-agent hierarchical reinforcement learning approach. All agents learn to walk, navigate, and resolve behavioural constraints within the same environment amongst each other.

Finally, generality is often the product of very large simulation or datasets, requiring similarly large compute resources. It is possible to define a type of combinatorial curriculum learning where increasing types and numbers of agents afford data-efficient learning of more general policies which generalize better to different scenarios.

## V. Conclusion

This work argues looking at synthetic crowds critically, toward developing much more sophisticated models for future use in safety critical and human predictive applications such as inclusive built environment design. The avenues for solutions to the three issues of data, diversity, and generality in data-driven DRL methods form the basis of a research program seeking to create representative synthetic crowds models for inclusive environment design, policy development, and safety-critical analysis.

## References


[1] M. A. Jackson, "Models of disability and human rights: Informing the improvement of built environment accessibility for people with disability at neighborhood scale?" *Laws*, vol. 7, no. 1, p. 10, 2018.

[2] R. Garland-Thomson, "Feminist disability studies," *Signs: Journal of Women in Culture and Society*, vol. 30, no. 2, pp. 1557–1587, 2005.

[3] S. Wendell, *The Rejected Body: Feminist Philosophical Reflections on Disability*. Routledge, 1996.

[4] A. Renalds, T. H. Smith, and P. J. Hale, "A systematic review of built environment and health," *Family & community health*, vol. 33, no. 1, pp. 68–78, 2010.

[5] P. Clarke, J. A. Ailshire, M. Bader, J. D. Morenoff, and J. S. House, "Mobility disability and the urban built environment," *American journal of epidemiology*, vol. 168, no. 5, pp. 506–513, 2008.

[6] P. Clarke and L. K. George, "The role of the built environment in the disablement process," *American Journal of Public Health*, vol. 95, no. 11, pp. 1933–1939, 2005.

[7] G. Berseth, M. Kapadia, and P. Faloutsos, "Robust space-time footsteps for agent-based steering," *Computer Animation and Virtual Worlds*, 2015.

[8] S. Singh, M. Kapadia, G. Reinman, and P. Faloutsos, "Footstep navigation for dynamic crowds," *Computer Animation and Virtual Worlds*, vol. 22, no. 2-3, pp. 151–158, 2011.

[9] G. Berseth, M. Kapadia, B. Haworth, and P. Faloutsos, "SteerFit : Automated Parameter Fitting for Steering Algorithms," in *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. Eurographics Association, 2014, pp. 1–10.

[10] J. M. Finley and A. J. Bastian, "Associations between foot placement asymmetries and metabolic cost of transport in hemiparetic gait," *Neurorehabilitation and neural repair*, vol. 31, no. 2, pp. 168–177, 2017.

[11] M. M. Platts, D. Rafferty, and L. Paul, "Metabolic cost of over ground gait in younger stroke patients and healthy controls." *Medicine and science in sports and exercise*, vol. 38, no. 6, pp. 1041–1046, 2006.

[12] K. M. Michael, J. K. Allen, and R. F. Macko, "Reduced ambulatory activity after stroke: the role of balance, gait, and cardiovascular fitness," *Archives of physical medicine and rehabilitation*, vol. 86, no. 8, pp. 1552–1556, 2005.

[13] O. Nachum, H. Tang, X. Lu, S. Gu, H. Lee, and S. Levine, "Why does hierarchy (sometimes) work so well in reinforcement learning?"

[14] B. Haworth, G. Berseth, S. Moon, P. Faloutsos, and M. Kapadia, "Deep integration of physical humanoid control and crowd navigation," in *Motion, Interaction and Games*, ser. MIG '20. New York, NY, USA: Association for Computing Machinery, 2020. [Online]. Available: https://doi.org/10.1145/3424636.3426894

[15] X. B. Peng, G. Berseth, K. Yin, and M. van de Panne, "Deeploco: Dynamic locomotion skills using hierarchical deep reinforcement learning," *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2017)*, vol. 36, no. 4, 2017.

[16] W. Yu, V. C. Kumar, G. Turk, and C. K. Liu, "Sim-to-real transfer for biped locomotion," *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nov 2019. [Online]. Available: http://dx.doi.org/10.1109/IROS40897.2019.8968053