Modelling and simulation of human behaviour for human-centric digital twins

Victor Lavairye[a,*], **Chris Reinke**[a], **Florian Noyrit**[a] **and François Terrier**[a]

[a]Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France
ORCiD ID: Victor Lavairye https://orcid.org/0009-0001-5650-609X,
Chris Reinke https://orcid.org/0000-0003-2368-521X, Florian Noyrit https://orcid.org/0000-0001-5947-7258,
François Terrier https://orcid.org/0000-0002-6871-9818

Abstract. Digital twins are increasingly adopted to enhance the design, operation, and optimisation of complex systems, yet their effectiveness often relies on the accurate representation of human behaviour. Accurately capturing human behaviour remains a key challenge, particularly in configurations where no empirical data exists. Here, we present a method that produces plausible human-like behaviours in previously unobserved configurations. The method combines simulation-based generation of candidate behaviours with a data-driven selection step that ranks candidates by similarity to human behaviours observed in previous configurations. We apply this approach to a large-scale assembly task data set, Assembly101, with ongoing work focused on its development and evaluation.

1 Introduction

Digital twins (DT) are virtual replicas of physical systems that mirror real-world behaviour in real time for simulation, analysis, and optimisation [7]. They can be used to explore, via simulations, how the system would respond to hypothetical changes, i.e., what-if simulations, prior to implementing them in the physical system.

However, providing accurate insights is particularly challenging in systems that closely interact with humans — often referred to as *human-centric DTs* — because it requires accounting for the complexity and richness of human *behaviour*. This necessitates modelling human behaviour in novel configurations, even when no empirical data is available. This leads us to our research question: **How can plausible human behaviour be generated in a new DT configuration for which there is no data?**

In this work, we focus on the problem of generating human-like behaviour for assembly tasks. Assembly tasks are error-prone, *goal-oriented* procedural processes in which goals and actions appear at multiple levels of abstraction.

2 Overall Approach

Task. Given a simulation environment for an assembly task and human behavioural data from previous tasks, our goal is to generate plausible human-like behaviours for a new task for which there is no data.

Approach. The approach follows a two-step procedure:

1. **Generation**: produce a set of diverse behaviours using the simulator of the new task.
2. **Selection**: filter this set to retain behaviours most similar to human ones observed in other tasks.

The first step is unbiased and seeks coverage, while the second step injects human-likeness.

3 Use case: Assembly101

Data set. Assembly101 [1, 8] is a large-scale data set of humans assembling 15 types of vehicle toys (e.g., cars and trucks), comprising over 300 labelled assembly video sequences. Each sequence provides temporal annotations of actions, labelled with a verb (attach or detach) and the involved toy components. Actions that must later be corrected to complete the assembly are explicitly labelled as *erroneous*, capturing instances of human mistakes and making the data set suitable for studying human behaviours.

World model. As no simulator currently exists for Assembly101, we model the assembly of each toy as $\mathcal{M} = (C, \mathcal{R}, \mathcal{P})$ (Figure 1, left), where C is the set of components, $\mathcal{R}$ is an irreflexive, symmetric, and acyclic relation over C encoding possible attachments, and $\mathcal{P}$ is an order relation over C capturing temporal constraints: if $(c, c') \in \mathcal{P}$, then c must be attached before c', as their shapes make the reverse order physically impossible.

World states are spanning subgraphs of $(C, \mathcal{R})$, denoted w. Actions correspond to edge additions or removals. Given a state w, the feasibility of an action a is expressed by the formula $\mathcal{M}, w \models a$. This formalism induces a *labelled transition system* (LTS), used as a simulator.

Behaviours. According to the world model described above, a *behaviour* is a finite sequence of actions $\langle a_1; \ldots; a_n \rangle \in A^n$ (see Figure 1, right for an example).

Problem setup. Let z be the new toy to be assembled, and let $\mathcal{T}$ denote the set of previously assembled toys with available human data. The objective is to generate a set of plausible human-like behaviours for the new toy, $\hat{B}^z$, using its world model $\mathcal{M}_z$ and the observed human behaviours from $\{B^t\}_{t \in \mathcal{T}}$.

3.1 Generation of diverse behaviours

Challenge. Our aim is to generate a set of behaviours that includes human-like behaviours. However, while the number of states grows

* Corresponding Author. Email: victor.lavairye@cea.fr

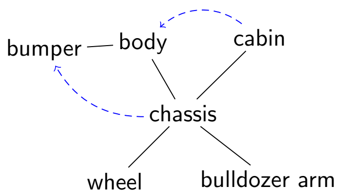

b =⟨attach(chassis, wheel);
attach(body, chassis);
attach(bulldozer arm, chassis);
detach(body, chassis);
attach(cabin, chassis);
attach(body, chassis);
attach(body, bumper)⟩

Figure 1: Left: Graphical representation of a world model $\mathcal{M}$, where nodes are components (C), solid edges are possible attachments ($\mathcal{R}$), and dashed curved arrows indicates temporal constraints ($\mathcal{P}$). Right: Example of a possible behaviour b with respect to this model.

exponentially with the number of components ($2^{|C|-1}$), the number of possible behaviours grows factorially, making naive trajectory space exploration infeasible. In addition, the generation process must be computationally efficient.
Techniques. To address this, we investigate search algorithms that efficiently explore trajectory space by focusing on diversity:

- *Intrinsically Motivated Goal Exploration* [5]: is a learning paradigm in which agents autonomously generate and pursue goals, guided by intrinsic rewards such as novelty, curiosity, or progress in learning. This enables efficient and open-ended skill acquisition without relying on predefined external objectives, often leading to more general behaviours.
- *Novelty search* [2, 3]: is an *evolutionary algorithm* that rewards behavioural novelty instead of progress toward a predefined objective, thus encouraging the discovery of behaviours that differ significantly from those previously encountered.
- *Quality-Diversity* [4]: aims to discover a diverse set of solutions that are not only behaviourally distinct, but also exhibit high performance according to a predefined quality metric, each representing the best candidate within a specific behavioural niche.

In our setup, these are not learning agents but sampling strategies applied to the LTS induced by $\mathcal{M}_z$, aiming purely at diversity.

3.2 *Selection of human-like behaviours*

Challenge. Once candidate behaviours are generated, the central question is how to identify those that resemble human behaviour observed in previous tasks. The key difficulty is defining a distance measure between behaviours that remains meaningful across assemblies of toys with different components and structures.
Distance measure between behaviours. Behaviours can be compared by the *Levenshtein distance*, which measures the cost required to transform one behaviour into another by means of edit operations on individual actions: *substitution*, *insertion*, and *deletion*. While it captures the syntactic form of behaviours (order and length), it ignores functional similarities across components.

In Assembly101, components with different names can serve an equivalent role. Similarly, identically named components may perform different roles. To account for this, we extend the Levenshtein distance by weighting edit costs using a *semantic proximity* between components. This proximity is computed as the cosine distance between embeddings from a *vision-language model* (e.g., CLIP [6]), so that substituting semantically related parts — for example, "bulldozer arm" for "excavator arm" — incurs a lower distance than substituting semantically unrelated components. The resulting distance provides a more meaningful measure of behavioural closeness across tasks.
Selection of human-like behaviours. For each generated candidate behaviour b, we compute

$$L(b) = \min_{t \in \mathcal{T}, b' \in B^t} d(b, b').$$

The candidates with the smallest values are retained as the human-like behaviours for the assembly of the toy z.

Optionally, to ensure robustness and avoid outliers, selection can be restricted to candidates that belong to a dense cluster in the behaviour space, ensuring the final set $\hat{B}^z$ fosters coherent human-like strategies rather than isolated sequences.

3.3 *Evaluation of the approach*

Human-like performance. The performance of the approach is evaluated by comparing the selected behaviours with the ground-truth human behaviours for the assembly of z. This comparison is carried out using the same distance described in Section 3.2.
Data efficiency. Since the approach relies heavily on data, it is important to quantify how performance varies with data availability. To this end, we measure the performance of the approach using subsets of $\mathcal{T}$ of decreasing size, removing one toy at a time, and compare the results across these different subsets.
Computation efficiency. What-if simulations must be generated within a sufficiently short time frame to provide feedback without incurring substantial delays. Consequently, to ensure that the approach is applicable in the context of digital twins, a study of its computational efficiency will be carried out.

4 Conclusion

This work addresses the challenge of generating plausible human behaviour in digital twins for unseen configurations. We have introduced a symbolic world model grounded in component structure and constraints, and developed a two-stage pipeline that first generates diverse behaviours and then filters them using refined metrics for human-likeness.

Ongoing work includes implementing search-based generation, developing semantic-aware distance measures, and running evaluations on Assembly101.

Acknowledgements

This work is carried out under the supervision of François Terrier (CEA-List) and the advisoryship of Florian Noyrit (CEA-List). I would also like to thank Chris Reinke (CEA-List) for his support and advice.

References

[1] Guodong Ding, Fadime Sener, Shugao Ma, and Angela Yao, 'Spatial and temporal beliefs for mistake detection in assembly tasks', *Computer Vision and Image Understanding*, **254**, 104338, (March 2025). Publisher: Elsevier BV.

[2] Joel Lehman and Kenneth O. Stanley, 'Improving evolvability through novelty search and self-adaptation', in *2011 IEEE Congress of Evolutionary Computation (CEC)*, pp. 2693–2700, (2011).

[3] Joel Lehman and Kenneth O. Stanley, *Novelty Search and the Problem with Objectives*, 37–56, Springer New York, New York, NY, 2011.

[4] Justin K. Pugh, Lisa B. Soros, and Kenneth O. Stanley, 'Quality diversity: A new frontier for evolutionary computation', *Frontiers in Robotics and AI*, **Volume 3 - 2016**, (2016).

[5] Alexandre Péré, Sébastien Forestier, Olivier Sigaud, and Pierre-Yves Oudeyer. Unsupervised learning of goal spaces for intrinsically motivated goal exploration, 2018.

[6] Alec Radford, Jong Wook Kim, Christine Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pam Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever, 'Learning transferable visual models from natural language supervision', *arXiv preprint arXiv:2103.00020*, (2021).

[7] Benjamin Schleich, Nabil Anwer, Luc Mathieu, and Sandro Wartzack, 'Shaping the digital twin for design and production engineering', *CIRP Annals*, **66**(1), 141–144, (2017).

[8] Fadime Sener, Dibyadip Chatterjee, Daniel Shelepov, Kun He, Dipika Singhania, Robert Wang, and Angela Yao. Assembly101: A large-scale multi-view video dataset for understanding procedural activities, 2022.