URBAN ANALYTICS AND GENERATIVE DEEP LEARNING FOR CONTEXT RESPONSIVE DESIGN IN DIGITAL TWINS

A Singapore Study

IBRAHIM NAZIM[1] and SAM CONRAD JOYCE[2]
[1,2]*Singapore University of Technology and Design*
[1]*ibrahimnazim@hotmail.com, https://orcid.org/0000-0003-1971-2254*
[2]*sam_joyce@sutd.edu.sg, https://orcid.org/0000-0001-6637-4193*

Abstract. City Digital Twins (CDT) can play a pivotal role in consolidating and visualising complex urban big data, promoting rapid and informed decision making in contemporary cities. Beyond rich contextual data, these tools offer features like interactivity, 3D models and data visualisation, making them ideal for urban planning and design explorations. However current CDT implementations primarily focus on data visualisation and lack any robust design support. Simultaneously, generative urban prototyping occurs in specialised tools, detached from this rich contextual data. This study, on the Virtual Singapore (VSg) CDT platform, explores how the platform's existing data, interactivity and 3D visualisation capabilities can facilitate rapid prototyping through generative machine learning techniques trained on the city's unique planning texture; and discusses the challenges and limitations of the platform in supporting the development of such tools, along with potential improvements. The study aims to advance CDTs beyond static data consumption and visualisation towards generative tools for urban planning and decision-making processes.

Keywords. City Digital Twins, Virtual Singapore, Generative design, Planning Support, Smart City.

1. Introduction

Contemporary cities face complex social, environmental and economic challenges that demand rapid decision making and adaptive planning. In this dynamic context, the availability of planning support tools, ensuring timely access to aggregated data and analytics, is crucial for effective decision making. The emergence of Big Data and Smart City strategies, including the internet of things (IoT), has facilitated the generation and collection of substantial data on the city and its residents. Meanwhile, advancements in data analytics and ML has empowered planners to efficiently explore and analyse this data to extract valuable insights. By consolidating these diverse data sources, planners can derive location specific analysis and receive real-time feedback on factors such as walkability of urban areas (Cheng et al., 2022).

ACCELERATED DESIGN, Proceedings of the 29th International Conference of the Association for Computer-Aided Architectural Design Research in Asia (CAADRIA) 2024, Volume 2, 495-504.

These advancements have given rise to City Digital Twins (CDT), which are virtual representations of cities driven by data and 3D models (Schrotter and Hürzeler, 2020). These models consolidate data from diverse sources to create detailed representations of cities, encompassing layers of information on the physical environment, including buildings, infrastructure etc alongside, population demographics, social activity, and real-time sensor data. Potentially anything that is spatial and relevant to planning and managing a city can be added. Pilot projects for smart CDT are underway in cities like Singapore, Glasgow, Zurich, Helsinki, and Dublin (Caprari, 2022). These initiatives, leveraging simulation, visualisation, and scenario modelling, aim to empower government agencies, urban planners, and researchers to collaborate and make informed policy decisions (Papyshev and Yarime, 2021).

Simultaneously, recent advances in computational design and generative AI has significantly enhanced our capacity to rapidly synthesise design solutions. Specifically deep learning approaches, trained on large volumes of existing planning data, can internalise existing design rules weather explicit or informal via precedent, unlike parametric or algorithmic approaches before this. The benefits of generating and evaluating multiple design solutions in the early stages of design have been well documented, especially when integrated with environmental and urban performance metrics (Düring et al., 2020). However, these explorations often take place within specialised design tools, disconnected from the rich contextual data available. Integrating the two involve building custom tools and workflows that require specialist knowledge and manual consolidation of data from various sources. Although there are some new and promising tools aiming to bridge this gap, such as Digital Blue Foam, Giraffe, TestFit and Delve, they currently focus on specific aspects of the context and therefore lack the data richness and comprehensiveness available in CDTs.

In contrast, CDTs offer rich contextual information, coupled with 3D visualisation, interactivity and collaboration capabilities, all of which enhance effective design exploration. However, they conventionally function primarily as data consumers and aggregators, As a result lacking any significant design capabilities, and thus limiting their direct participation in the design process. This study, conducted within the Virtual Singapore (VSg) CDT platform, investigates the potential introduction of generative design capabilities in CDT environments. The aim is to leverage their rich contextual information and visualisation capabilities for an integrated urban prototyping and analysis process. Specifically seeking an approach that is tailored for the typical use of CDT; that of high-level strategic planning which prioritises fast low effort exploration of options that are feasible and likely given existing rules and norms.

2. Background

2.1. VIRTUAL SINGAPORE

VSg is a collaborative 3D city model of Singapore, established through a partnership involving the National Research Foundation (NRF) and various government agencies, with the primary aim of facilitating virtual testing and experimentation for planning and decision-making (SLA, 2014). Hosted on Dassault Systèmes' 3DExperience service, it is an interactive web platform meaning unlike many other GIS like tools it is accessible to all with an internet browser, allowing the import, export, and visualisation of various datasets at different levels of detail (LOD). It contains building information of the whole island in LOD2, utilising 2014 data. Specific areas have been refined to LOD2.5 and LOD3, incorporating detailed models of streets, trees and urban fixtures (Papyshev and Yarime, 2021). However, due to safety considerations, access to much of this data is restricted, available solely to their owners and specific user groups (Reuters, 2018).

Beyond its built-in functionalities for visualisation, collaboration and analysis, the VSg platform enables development of custom tools, as specialised widgets, utilising their API to communicate with its data and model. This allows for significant system extensibility by linking to external systems of data, analysis as well as simulation. Researchers have harnessed this capability to build their own tools for analysis and visualisation, such as integrated climatic simulation visualisation (Ignatius et al., 2021). However, there are no studies exploring VSg for design generation, and few exploring the use of web-based CDT systems for design in general. We believe one of the reasons for this is the general incompatibility between the complex interaction needed in typical designs systems compared to the more passive data exploration and consumption CDTs are orientated to provide. However, the solution is not to overload an already complex CDT with CAD like functionality but rather conceive of a more appropriate design exploration approach. To do this we believe recent ML driven generative design offers a compelling alternative.

2.2. GENERATIVE URBAN DESIGN

Most generative urban design approaches currently rely on “procedural” or “parametric” modelling techniques utilising specialised software like Esri CityEngine and Rhino3D's Grasshopper. In such cases, predefined rules and operations are employed to generate specific building typologies and other urban elements (Stouffs and Janssen, 2017). The advantage is that once the initial rules and model has been set up, exploring options is quick and automated, and easily paired with analysis and optimisation tools (Koenig et al., 2020). However, these approaches typically rely on very basic rules, resulting in simplistic forms that lack the sophisticated geometries as well as responding to the often complex interacting legal rules and subtle implicit sensibilities of urban design that real buildings adhere too. While solutions exist to enhance the diversity of parametric models (Harding et al., 2013, Ibrahim and Joyce, 2020) capturing the complexity of existing urban contexts through rule-based techniques remain challenging.

In contrast, ML techniques such as Generative Adversarial Networks (GAN) can learn patterns from training images and generate new images consistent with the training data. Previous studies have utilised GAN models to predict urban features such as building footprints and road networks consistent with existing patterns (Wu and Biljecki, 2022, Boim et al., 2022) and for predictive planning and design (Chaillou, 2020, Tian 2020). While there are limitations to their capabilities, the resolution proves effective for early-stage urban design explorations (Joyce and Nazim, 2021). Additionally, once trained, these models can be easily integrated into a webpage, allowing the generation of new scenarios based on image or interactive sketches as inputs. This enhances compatibility with web-based CDTs, and improves collaboration and accessibility for planners without specialised computational tools or expertise.

3. Methodology

The study focussed on developing an urban planning process for designing new residential developments based on existing housing typologies of Singapore that could be a good fit with the process flow and interface paradigm of the VSg system. To facilitate this, two custom widgets were implemented. Firstly, a visual analysis support tool was introduced, offering a neighbourhood score based on surrounding amenities. Secondly, a GAN-based generative design tool was incorporated, allowing the interactive generation of buildings and trees. Both widgets utilised the platform API for model interaction, retrieving contextual data and geometry, as well as importing and visualising the results. Due to limited access to existing platform data, the current implementations relied on importing publicly available data from data.gov.sg, sourced from official sources. This included plots and land uses from URA masterplan data, and locations for amenities and other destinations from various relevant agencies.

3.1. VISUAL ANALYSIS SUPPORT TOOL

3.1.1. Data and Neighbourhood Score

An overall neighbourhood score was implemented based on the distance and diversity of land uses and amenities within a given walkshed of a selected point. For land uses, distances were measured to the centroid of the polygon. For each category of land use, POI or transit data, the individual distances were aggregated to get the average and minimum distance (minDist). Distance score (dScore) for each category (c) is then calculated as:

dScore(c) = 1- (minDist(c)/walkshed radius)

The average dScores for all destinations was considered as the Walkability score while the average for transit nodes and bus stops was considered as the Transit score. Diversity scores were calculated using Shannon-Weaver diversity index (Yoshimura et al. 2022):

$H = -\sum[(pi) * \log(pi)]$, where $pi = n/N$

n = individuals of a given species, N = total individuals in the community

A Land-use diversity score based on proportion of areas and a Destination diversity score based on proportion of counts of destinations was calculated. The overall neighbourhood score was then calculated as the average of the Walkability, Transit, Land-use diversity and Destination diversity scores.

3.1.2. Implementation

The implemented widget allows users to create a walkshed by clicking on a point or a plot on the model and setting a walkshed radius using a slider input. Any data sources to be used for land uses, destinations and transit scoring can be added from available data layers using the relevant data selection tabs in the widget. For easier decision making, the analysis results and metrics were presented in various levels of details as a dashboard (Figure 1). The dashboard presents high-level metrics such as overall neighbourhood score as well as detailed analysis for each of the data types. It allows users to interactively explore the data based on counts of locations, minimum or average distance and the total area for land-uses.

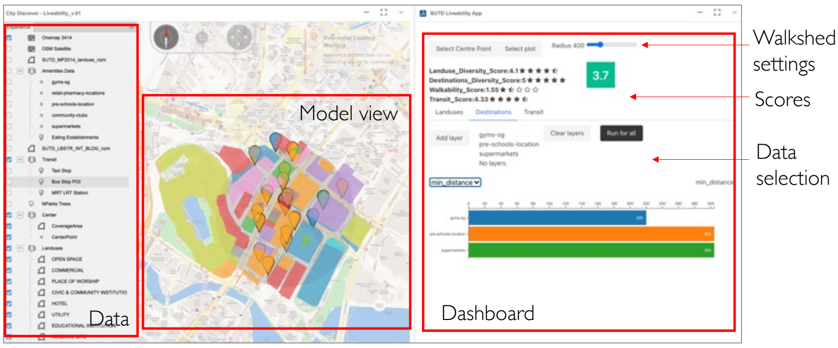

Figure 1 Elements of the visual analysis support widget within VSg context

3.2. GENERATIVE DESIGN SUPPORT TOOL

3.2.1. Data and Training

The primary objective of this widget was to provide a high-level low-complexity design tool that could generate likely buildings and vegetation placement. Rather than a complex direct CAD like interface that VSg is not developed to easily support, an intelligent system that emulates typical commercial and public housing designs in Singapore whilst allowing for adaption to local context and a new variation from one click. Thus, something that generated based on existing data which was known to be compliant with or at least allowed by existing planners and rules.

GIS datasets of plots and buildings from LTA and tree locations from NParks were used as primary data sources for training the GAN models. A workflow to generate stylised image tiles from GIS data was implemented in Python (Figure 2). In addition to creating specifically stylised tiles, the workflow allowed filtering sites based on queries such as land use, type of building, land area etc.

Figure 2 Samples of Singapore housing typologies (Left) and training data (Right)

Pairs of training images for inputs and outputs were generated and split into sets, training, testing, and validation. Different GAN models were trained on this data using the Pix2Pix implementation (Isola et al. 2018). This is a conditional GAN, where a target image is generated, conditional on a given input image. It aims for images that are both plausible within the target domain, and plausible translations of the input image. This method was used to train and tune multiple models, of which two were chosen and integrated into the final widget (Figure 3).

Model 1 takes in an input image of an empty site (magenta) and its surrounding context (cyan) and roads (white) and predicts the building footprint for the given site. The model was trained on a subset of the data, consisting of only residential plots greater than 10,000 sqm. Model 2 takes in an input image of roads (white), buildings (black), plots (grey) and open areas (green) and predicts locations for trees. This model was trained by combining the SLA and NParks tree data set. Other models developed but not implemented include a model for predicting missing road configurations for a given area and a model to predict building contexts just from road networks, but these had unacceptably low accuracies, and generally created incoherent designs.

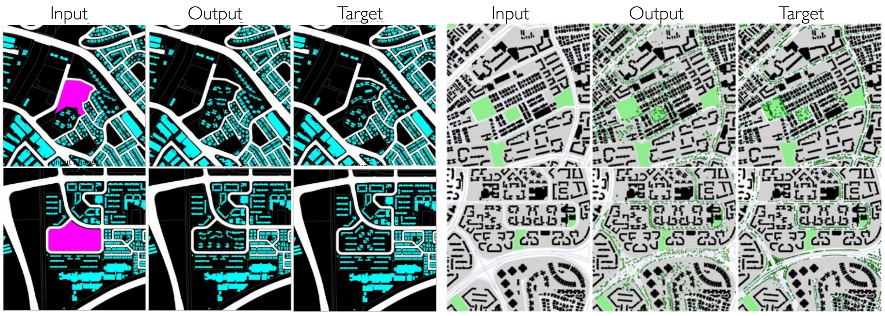

Figure 3 Examples of input, original, and generated images for Model 1 (Left) and Model 2 (Right)

3.2.2. Implementation

The trained models were hosted online and made available through a simple web page that is served as part of the widget on the VSg platform and integrated with the platform API to generate input images directly from the map. The platform API and functionality is utilised for selecting the data required, such as the target site, road configuration, and building context, to generate the input tile as well as visualising the output 3D geometry. To support an interactive exploration process, several features were implemented (Figure 4).

The model view represents the tile boundary and utilises built-in 3D viewer to visualise 3D geometries of prediction building footprints as well as trees. The actions bar contains a drop down to select the prediction model to be used as well as buttons to all the key operations of the process. This includes selection of site, adding plot and building geometries for the input maps as well as generating and updating the input tile. The generative interface provides tools to interactively modify the input tile by painting over it. The stroke size and colour can be adjusted using the inputs provided. Compared to existing systems this is a simple interface, with a low number of parameters and simple site-based input. Instead allowing for user to simply regenerate to get new options or use the sketch like interface to sculpt the output.

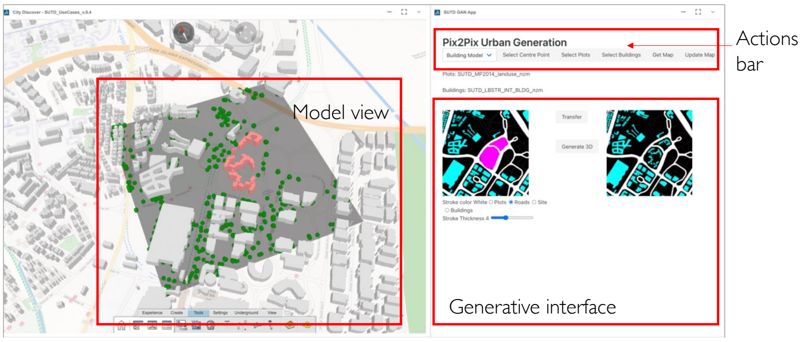

Figure 4 Features of the generative design widget within VSg context

4. Case Study

In this section, we demonstrate how the two widgets can be used in succession as part of an integrated urban planning process to prototype new residential developments. The analysis widget is used to identify potential development sites based on their neighbourhood scores and the generative design widget is used to generate footprints and prototype design options for the top performing neighbourhoods. The generated options and analysis are then saved and shared with decision makers via built-in collaborative functionalities of the platform.

Five key regions in Singapore were chosen as potential areas based on main transit access; Holland Village, Orchard, Tanjong Pagar, Punggol, Jurong East. For each, neighbourhood scores were generated using the analysis widget within a radius of 800m around the station, which is considered an ideal distance when planning around transit area (Gori et al. 2014). Through comparative analysis of the five regions, the top 3 sites were selected based simply on the overall neighbourhood score. For the selected regions, potential residential plots were identified by visualising their land-use attribute using built in visualisation functions. For each site, multiple design options were generated using the sketching feature and visualised in 3D by applying heights to the generated shapes. As no evaluation metrics were implemented at this stage, final building forms were decided purely based on satisfaction of the designer (Figure 5).

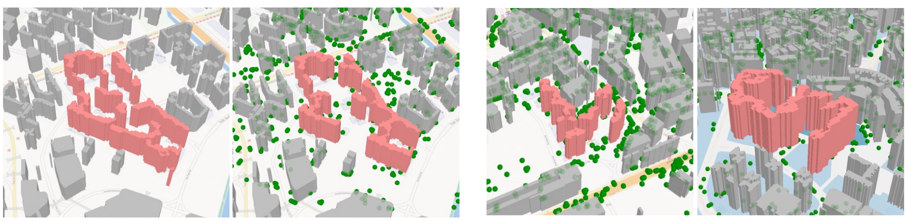

Figure 5 Two design options for one site (Left) and final designs for the other two top sites (Right)

The generated designs were saved in a shared workspace on the platform and bookmarked for easy navigation for other collaborators to explore. The same process was also applied iteratively to multiple sites in one of the neighbourhoods to envision a phased development. Here the initial plot decisions affected subsequent neighbouring plots and separate design exploration sessions generated varying results revealing an adaptive densification process (Figure 6).

Figure 6 Results of two separate iterative design generation sessions on one neighbourhood demonstrating adaptive densification in both examples.

5. Discussions and Conclusions

5.1. PLATFORM FEATURES AND LIMITATIONS

Despite privacy limitations in working with 3D datasets, meaning the study was in 2.5D, the platform’s web-based environment offered numerous features that greatly facilitated tool development. Since widgets function as standalone webpages hosted externally, dynamic interfaces and dashboards can be implemented and updated using established web technologies while seamlessly integrating with various APIs and analysis systems, like live traffic cameras and weather APIs. Crucially, the platform’s API enables interaction with model and data layers, allowing the leverage of platform data as input for analysis, as well as injecting new visualisations and geometry into the VSg environment as output. The cloud-based storage of these elements enhances accessibility, enabling easy option capture, tool deployment at scale, and sharing among users through the built-in social network platform for collaborative feedback.

Simultaneously, the whole platform being hosted on the web led to longer loading times in the browser, especially with models containing multiple datasets. Requiring a persistent server rather than an on demand client side software tool. Despite using techniques like tiling to improve performance, the platform often strained browser resources. Furthermore, as the platform primarily focuses on collaboration and visualisation and keeping much of this data server side, rather than external data transferers, there are limits to built-in analytical functions and prevents easy use of eternal tools or simulations. While it supports basic measurements, area calculations and visibility analysis on imported models, these are not available for generated geometries and therefore need to be implemented using third party libraries. Providing these along with other essential urban analytics such as 3D isovist and isochrone analysis via the API can significantly enhance the development of more robust tools.

5.2. APPLICATIONS AND FUTURE WORK

The study showcased the integration of generative design and analysis workflows for rapid urban prototyping in web-based city digital twin platforms like VSg. The analysis widget derived location-specific metrics and heuristics from multiple data sources, simplifying planners' analysis tasks. Similarly, the generative design application demonstrated context-specific ML models to rapidly generate new designs consistent with surrounding features, offering multiple options for the same site using integrated sketch tools. However, both systems have room for improvement. Feedback from 15 participants, from the research community, who tested these tools during two online workshops, emphasised the value of predictive and auto-suggestive design features for ideation. However, improvements are required for better prediction control and integration of analysis and simulation for goal-oriented design.

To enhance the analysis process, additional scores can be easily implemented by leveraging different forms of datasets including business and employment data, demographic information, land values as well as mobility data. These can aid in assessing infrastructure and resources accessibility, identifying potential development areas or gain insights into how people move and use spaces in the city. Integrating and linking these diverse types of data allows participation of different stakeholders and can help to promote a shared practice towards more sustainable environments and cities (Vite et. Al, 2021). To ensure effective analysis support with data, it is crucial to develop goal-specific metrics and heuristics along with tailored analysis and visualisations. Furthermore, integrating post-generation analysis allows for the evaluation of potential implications arising from the generated scenarios. This enables planners to quickly assess various options for goal-oriented design, ranging from simple metrics like GFA, density, cost estimations and material usage to more complex performance simulations. While traditional simulations may be resource-intensive, employing a GAN-based approach to train proxy models can expedite environmental performance predictions, enhancing quick exploration tools like VSg (Yousif and Bolojan, 2021). While GAN models might not match the accuracy of simulations, they offer a swift alternative for early-stage evaluations.

As the current generative model is a conditional GAN, it inherently offers more control compared to traditional GANs by incorporating input labels as images. This can be enhanced by further refining input labels to capture additional information, such as plot density or building heights. These can be represented as alpha values or encoded as supplementary channels alongside the existing RGB channels. Additionally, extracting learned features from the models can enhance design exploration by analysing their latent space (Härkönen et. Al, 2020). Finally, the workflow can be optimised for design exploration by assembling and stacking multiple modular GAN models into larger workflows, each representing different stages of the design process (Chaillou 2020). Expanding beyond the current models for initial footprint and vegetation prediction, this approach could encompass internal layouts, external circulation as well as facades and landscaping.

Acknowledgements

This study was supported by the National Research Foundation (NRF), GovTech, and Dassault Systèmes. We thank them for their support and guidance.

References

Boim, A., Dortheimer, J., & Sprecher, A. (2022). A Machine-Learning Approach to Urban Design Interventions In Non-Planned Settlements.

Caprari, G. (2022). Digital Twin for Urban Planning in the Green Deal Era: A State of the Art and Future Perspectives. Sustainability.

Chaillou, S. (2020). ArchiGAN: Artificial Intelligence x Architecture. In Yuan, et al. (Eds.), Architectural Intelligence: Selected Papers from CDRF 2019. Springer.

Cheng, C., Li, Y., Deshpande, R., Antonio, R., Chavan, T., Nisztuk, M., Subramanian, R., Weijenberg, C., & Patel, S. V. (2022). Realtime Urban Insights for Bottom-up 15-minute City Design.

Düring, S., Chronis, A., & Koenig, R. (2020). Optimizing Urban Systems: Integrated optimization of spatial configurations.

Geddie, J., & Aravindan, A. (2018, September 27). Virtual Singapore project could be test bed for planners—And plotters. Reuters. From bit.ly/3I5bJsD

Gori, S., Nigro, M., & Petrelli, M. (2014). Walkability Indicators for Pedestrian-Friendly Design. Transportation Research Record.

Harding, J., Joyce, S., Shepherd, P., & Williams, C. (2013). Thinking Topologically at Early Stage Parametric Design. In Hesselgren et al. (Eds.), AAG 2012. Springer. Vienna.

Härkönen, E., Hertzmann, A., Lehtinen, J., & Paris, S. (2020). GANSpace: Discovering Interpretable GAN Controls. Advances in Neural Information Processing Systems.

Ibrahim, N., & Joyce, S. C. (2019). User Directed Meta Parametric Design for Option Exploration. 39th ACADIA. Austin.

Ignatius, M., Wong, N. H., Martin, M., & Chen, S. (2019). Virtual Singapore integration with energy simulation and canopy modelling for climate assessment. IOP Conference Series: Earth and Environmental Science,

Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. (2018). Image-to-Image Translation with Conditional Adversarial Networks (arXiv:1611.07004). arXiv.

Joyce, S. C., & Nazim, I. (2021). Limits to Applied ML in Planning and Architecture.

Koenig, R., Miao, Y., Aichinger, A., Knecht, K., & Konieva, K. (2020). Integrating urban analysis, generative design, and evolutionary optimization for solving urban design problems. Environment and Planning B: Urban Analytics and City Science.

Papyshev, G., & Yarime, M. (2021). Exploring city digital twins as policy tools: A task-based approach to generating synthetic data on urban mobility. Data & Policy.

Schrotter, G., & Hürzeler, C. (2020). The Digital Twin of the City of Zurich for Urban Planning. Journal of Photogrammetry, Remote Sensing and Geoinformation Science.

SLA. (2014). Retrieved 17 December 2023, from https://bit.ly/3I6bX2O

Stouffs, R., & Janssen, P. (2017). A Rule-Based Generative Analysis Approach for Urban Planning. In J.-H. Lee (Ed.), Morphological Analysis of Cultural DNA. Springer.

Tian, R. (2020). Suggestive Site Planning with Conditional GAN and Urban GIS Data. Proceedings of the 2020 DigitalFUTURES.

Vite, C., Horvath, A.-S., Neff, G., & Møller, N. L. H. (2021). Bringing Human-Centredness to Technologies for Buildings: An agenda for linking new types of data to the challenge of sustainability. CHItaly 2021: 14th Biannual Conference of the Italian SIGCHI.

Wu, A. N., & Biljecki, F. (2022). GANmapper: Geographical data translation. International Journal of Geographical Information Science, 36(7), 1394–1422.

Yoshimura, Y., Kumakoshi, Y., Milardo, S., Santi, P., Arias, J. M., Koizumi, H., & Ratti, C. (2022). Revisiting Jane Jacobs: Quantifying urban diversity. Environment and Planning B: Urban Analytics and City Science.

Yousif, S., & Vermisso, E. (2022). Towards AI-Assisted Design Workflows for an Expanded Design Space.