Demi Fang, Sophia V. Kuhn, Walter Kaufmann, Michael A. Kraus, Caitlin Mueller

Quantifying the Influence of Continuous and Discrete Design Decisions Using Sensitivities

Abstract: Early-stage design decisions significantly impact the performance of architectural geometries and structures, especially in the context of high-performance buildings. However, most design problems involve a combination of continuous and discrete variables as decisions, making it challenging to understand their individual contributions to performance. For instance, tradeoffs between geometry (continuous) and material type (discrete) often lead to non-obvious performance behaviors, complicating the inclination to establish "rules of thumb" for low-carbon structural design. In response, we propose a new approach for assessing the influence of mixed variables on performance based on variable sensitivities, with a focus on architectural applications. A conditional variational autoencoder is trained to learn a continuous representation of variables and their relationship to performance, and gradients of the learned model are used to compute the relative influences of variables on performance. Multiple scales of insights are attainable: 1) locally at a design instance of interest, and 2) globally across the whole design space. The method is demonstrated on a case study of designing a structurally feasible gridshell for low embodied carbon.

1 Introduction

In the face of the climate crisis, designers are more motivated to produce high-performing, low-carbon designs. Computational approaches offer a promising way to do so, combining quantitative performance evaluation with generative techniques. Parametric modeling and optimization techniques currently available for performance-driven design are valuable, but two ongoing challenges persist: 1) the technical challenge of exploring or optimizing mixed-variable (i. e. both discrete and continuous design variables) design spaces, which are common in architectural design but understudied in computational methods; and 2) the tendency for computational tools to prioritize providing designers with high-performing designs, rather than on revealing generalizable knowledge and enriching human intuition about the design problem.

Recently emerging computational methods can be leveraged to overcome these challenges and support new paradigms of performance-driven design. This paper proposes and demonstrates one such method: conditional variational autoencoders (cVAE) are used to build a learned representation of a mixed variable design space and performance as a continuous knowledge landscape. The resulting sensitivities can be computed automatically, and subsequently leveraged to provide both local and

https://doi.org/10.1515/9783111162683-031

global insights that inform design decision-making. This approach addresses both aforementioned challenges: 1) it offers a smooth way to handle discrete variables in combination with continuous ones, and 2) it offers a technically rigorous, data-driven way to produce insights on the design space. The method is demonstrated on a gridshell design case study.

2 Background

2.1 Navigating mixed variable spaces in design

Discrete variables may pose a technical challenge to design space exploration. Existing methods for design space exploration of mixed variable design spaces, in order from least to most rigorous, include:

1. Relying on "rules of thumb" (e. g. "always use timber where possible"), which lack specificity, carry cognitive bias, and are often too general.
2. Comparing just a few solutions across discrete choices, which risks omitting high-performing regions in the design space.
3. Evolutionary methods, which iterate through the design space but lack guarantees of convergence and a comprehensive understanding of influential decisions in addition to a significant computational effort.
4. Mixed integer programming (MIP), a well-established optimization method that provides a single optimal solution but lacks decision-making insights and typically requires commercial tools for implementation.

The proposed cVAE approach combines the benefits and overcomes the limitations of these existing strategies, providing a learned, continuous representation of the design space to reveal cognitive insights about design decision-making in the design problem.

2.2 Extracting design insights from sensitivity analysis

Various methods are available to quantify the influence of design decisions on performance. Classically, designers might attempt to understand design variable importance through Exploratory Data Analysis (EDA), for example by identifying steep slopes in objective-variable plots. This becomes impractical for high-dimensional design spaces.

Instead, we examine the opportunities offered by computing sensitivities, which can be offered locally or globally. Local sensitivity analysis examines the impact of input variable changes on the model output at a specific point and is widely used in various fields such as ecology and environmental modeling (Cho and Jung 2003). In contrast, global sensitivity analysis assesses these impacts (of inputs on outputs) across the entire design space and has been used in many areas, including engineering, physics, and economics (Saltelli et al. 2008).

Sensitivities are seldom used to provide insight in design decision-making, likely because they are challenging or expensive to compute directly for parametric design problems in which performance is computed via black box simulations or via hard-to-differentiate operations such as matrix inversions (e. g. FEM). Existing methods for approximating sensitivities include:

1. Finite differencing (e. g. Brown and Mueller 2018) which is limited to local recommendations (as opposed to global rules of thumb) in the design space due to linear approximations of performance function evaluations and are therefore both computationally expensive and inaccurate.
2. Automatic differentiation (AD) on a physics-based simulation (e. g. Hu et al. 2020). This is a promising future direction as exact analytical sensitivities are directly available. However, AD is not yet available for most performance function solvers, and sensitivities in the sense of derivatives are not defined for discrete variables.

This paper's cVAE approach employs AD to a learned continuous representation of the mixed variable design space to instantaneously obtain reasonably accurate sensitivity estimates once the model is trained.

2.3 Using generative models to produce design insights

Generative deep learning models are typically used in design for either optimization or "generative design" purposes (Regenwetter et al., 2022). Here we use a generative model for a different purpose: extracting design insights from a mixed variable design space.

Newer classes of generative models – e. g. conditional generative adversarial networks (cGAN) and conditional VAEs (cVAE) – improve upon conventional generative models by not only reducing dimensionality of high-dimensional design spaces but also incorporating performance as a condition; this enables the generation of high-quality designs that satisfy performance criteria (Danhaive and Mueller 2021). These developments are promising for performance-driven design applications, ensuring generated designs meet desired performance criteria during design space exploration. This work employs a specific cVAE architecture for co-learning a forward map between features and performances as well as a conditional distribution over the design space.

Despite cVAEs' potential to handle both challenges of mixed variable spaces and extracting sensitivities, few have demonstrated its usefulness in design decision-making. Balmer et al. (2022) recently present local and global sensitivities together with data visualizations for decision-making. We extend it to summarize decision-making insights as "influence metrics", a data-driven variation of the typical and manual "rules of thumb" familiar to designers.

3 Methods

The proposed method is summarized in Fig. 1 and in the following sections.

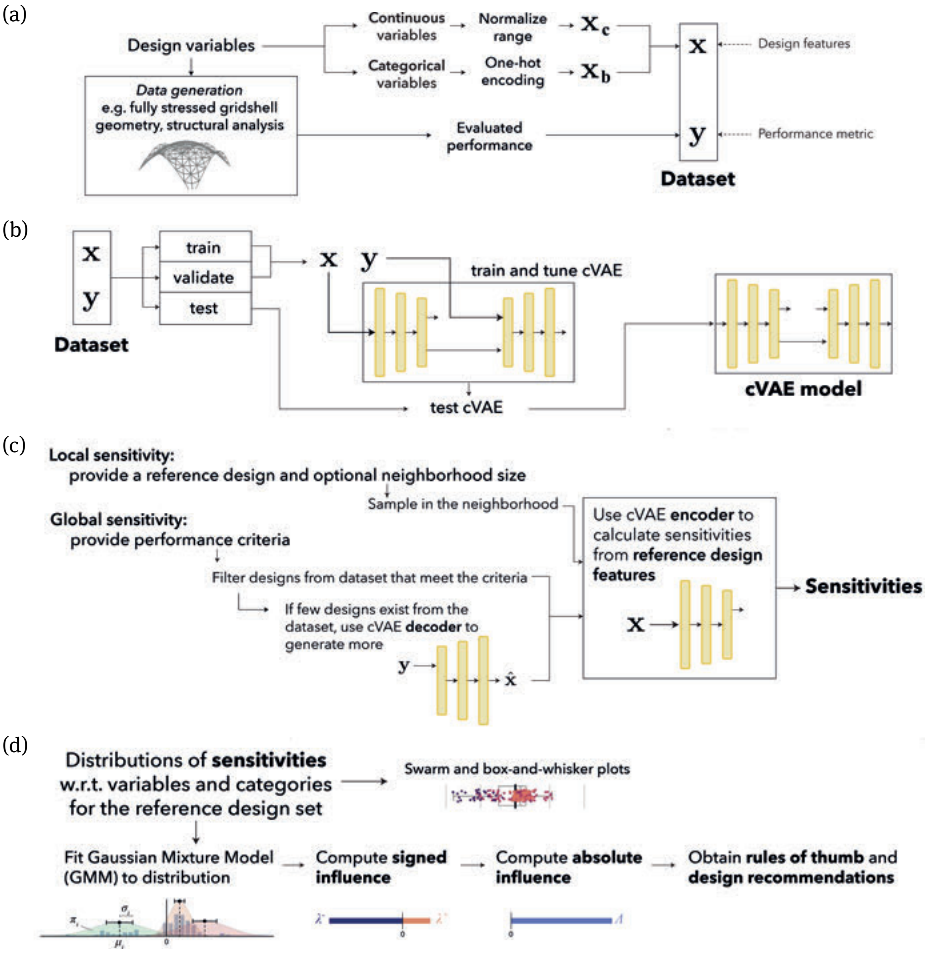

Fig. 1: Proposed method for extracting influence metrics from a mixed variable design space via cVAEs: (a) Data generation; (b) Build cVAE model; (c) Sensitivity analysis; (d) Influence metrics.

3.1 cVAE architecture

The basic deep learning model used in this study is the variation of the cVAE proposed by Balmer, Kuhn et al. 2022. In light of having to solve both a forward as well as an inverse problem, we forgo feeding the conditional **y** to the encoder and instead let it

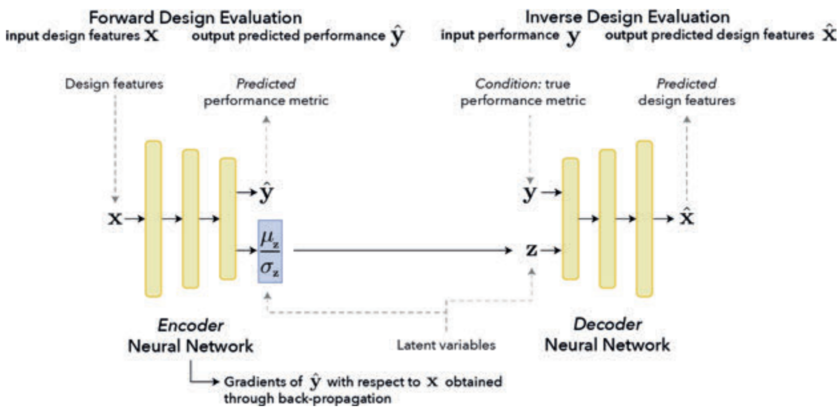

Fig. 2: The cVAE architecture enables solving both forward and inverse design problems. Adapted from Balmer et al. 2022.

predict the performance metrics together with a latent vector in two separate terminal nodes (Fig. 2).

Both the encoder and decoder consist of Multi-Layer Perceptron (MLP) Blocks with fully-connected layers, leaky-ReLU activation, and batch normalization. The total loss function L_{total} for training the cVAE is the sum of the reconstruction loss L_{des}, accuracy of predicted performance metrics L_{perf}, KL divergence L_{KL} , and decorrelation L_{cov} between $\mathbf{y}$ and $\mathbf{z}$, with hyperparameters w_i determining the contribution of each term:

$$L_{\text{total}} = w_1 \cdot L_{\text{des}}(\mathbf{x}, \hat{\mathbf{x}}) + w_2 \cdot L_{\text{perf}}(\mathbf{y}, \hat{\mathbf{y}}) + w_3 \cdot L_{\text{KL}}(\mathbf{z}) + w_4 \cdot L_{\text{cov}}(\mathbf{y}, \mathbf{z}) \tag{1}$$

Further details on background theory as well as the implementation may be found in Balmer, Kuhn et al. 2022.

3.2 Handling mixed variables: Pre-processing and obtaining sensitivities

The mixed variable space requires a new process for handling of continuous and discrete variables, summarized in Fig. 3. Discrete variables are common in AEC design and can be ordinal (e. g. integer) or categorical (unordered); this paper focuses on the latter.

Proper pre-processing is essential before passing design features to the cVAE to ensure comparable scales of sensitivities across variable bounds and units. Continuous variables are normalized to a range of [0, 1], and categorical variables are one-hot encoded. This variable pre-processing can be decoded later to retrieve the design variables in their original ranges and physical units.

For continuous variables, signs of sensitivities are correlated with an increase of the continuous variable value. For example, a positive gradient of embodied carbon

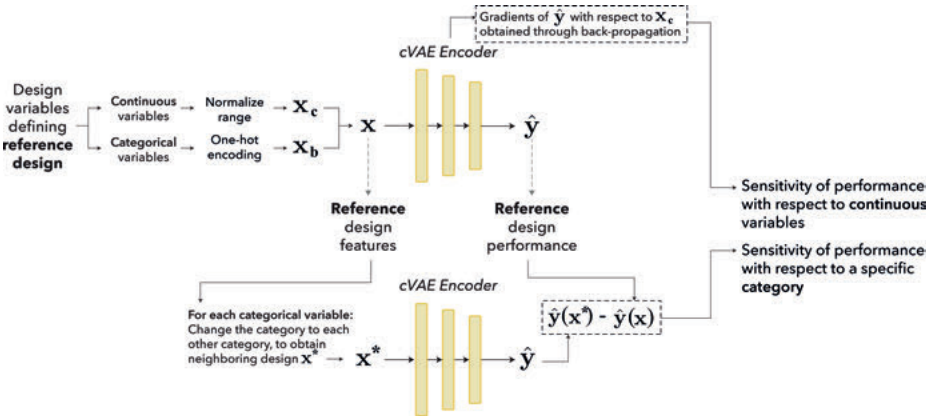

Fig. 3: Obtaining sensitivities of performance relative to continuous and categorical variables from a trained cVAE model at a reference design with design feature inputs **x**.

with respect to density indicates that an increase in density correlates with an increase in embodied carbon. This gradient can be obtained directly from the trained cVAE using AD on the encoder.

Categorical variables receive different sensitivity treatment. We extend the concept of sensitivity analysis for categorical variables, focusing on the impact of switching between specific categories. Rather than examining the overall sensitivity of embodied carbon to material choice, we propose a more relevant metric: the sensitivity of embodied carbon when transitioning from one category (e. g., steel) to another (e. g., timber). To achieve this, we introduce a separate evaluation method using gradients to assess performance with respect to categorical variables (Fig. 3).

3.3 Quantifying design variable influence

Presenting sensitivities in swarm and box-and-whisker plots (Balmer, Kuhn et al. 2022) can give the designer a sense of design variable importance.

Here, we utilize Gaussian mixture models (GMM) to algorithmically extract meaningful and rigorous insights from the data distributions, employing an established technique for approximating multimodal data as a learned weighted sum of normal distributions (Bischof and Kraus 2022).

The GMM returns means, variances, and relative weights of each distribution, which we utilize to compute the proposed influence and uncertainty metrics, summarized in Fig. 4. The absolute magnitude of these metrics is less meaningful than their relative magnitudes.

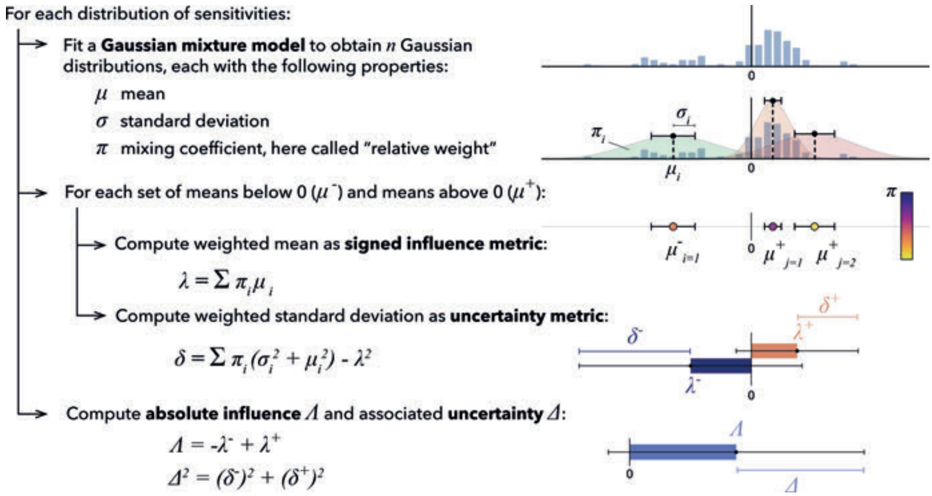

Fig. 4: Procedure for computing influence metric of a given distribution of sensitivities. For minimization performance metrics such as embodied carbon, negative influence metrics indicate decisions that improve performance, and vice-versa.

3.4 Example application

We evaluate the usefulness of our method by applying it to the design of a gridshell (Fang and Mueller 2021) containing a mixed variable design space (Tab. 1, Fig. 5). The effects of each variable on performance (embodied carbon) is somewhat intuitive, making it a good example application for evaluating the method. We employed standard material properties for steel (S235 per EN1993) and timber (C24 per EN1995), while for embodied carbon we use 1.23 kg CO2e/kg for recycled steel section and 0.51 kg CO2e/kg for timber glulam (no carbon storage) (Jones and Hammond 2019).

4 Results and reflections

The results of applying the proposed method on the example application are presented here. Figure 6 presents scatter plots of the objective-variable relationship, providing an overview of the generated design space. While this EDA aids in understanding performance trends, 2D projections can hinder insights, particularly for categorical variables where data points tend to overlap.

The trained cVAE model exhibits good predictive performance (coefficient of determination R^2 = 0.917, embodied carbon RMSE = 4786 kg CO2e). For brevity, training and model details are provided in a separate digital appendix (https://demifang.github.io/AAG2023/).

Tab. 1: Parameters of gridshell case study. Output is an embodied carbon estimate under fully stressed design assumptions.

Design variables	*Topology of gridshell:* Discrete {A, B, C, D, E, F, G} *Subdivisions:* Discrete, integer range varies by topology. (This variable is a designer input and a proxy for the "density" of the gridshell topology. A comparable metric, total length of all bars in the structure, is used in the sensitivity analysis to fairly compare "density" performance across topologies.) *Material (all gridshell bars):* Discrete {steel, timber}
Geometry generation	Topology is vertically projected onto a continuous shell of constant geometry.
Energy assumptions/ criteria	*Loading:* 3 kN/m^2 area load, applied as line loads – Full area – Quarter area for asymmetric loading *Strength:* designed to utilization ratio of 70 % *Serviceability:* maximum deflection of $L/200 = 6.7$ cm
Cross-section shapes	*Steel:* circular tube with diameter-to-thickness ratio of 20:1 – Series: outer diameters at 1-cm increments (4 cm to 37 cm) *Timber:* solid square section – Series: widths at 6-cm increments (2 cm to 68 cm)
Cross-section sizing	Sections are auto-sized to the smallest series section meeting the engineering design criteria. Auto-sizing occurs in two groups: edge bars and interior bars.
Performance	Embodied carbon (under fully stressed design assumptions), calculated by multiplying mass and material embodied carbon coefficient

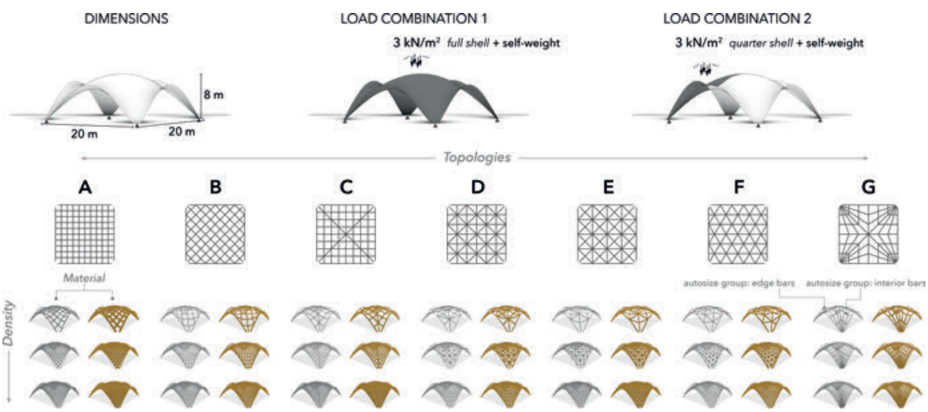

Fig. 5: Design space and engineering assumptions of gridshell example.

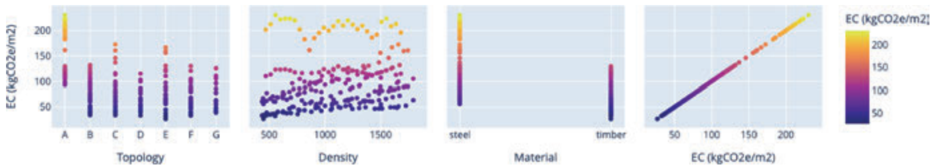

Fig. 6: Design space and engineering assumptions of gridshell example.

4.1 Local sensitivities: Action-oriented insights from a specific design

Local sensitivities assist decision-making when a designer is considering design decisions from a reference design. Rather than display the performance of adjacent designs and direct designers to select a specific design, sensitivities are presented as action-oriented insights, emphasizing the influence of each design decision.

We present local sensitivities and example interfaces for four reference designs in Fig. 7. For a chosen reference design, relative local sensitivities are shown as insights on best and worst design decisions. The results mostly agree with expert intuition (choosing topologies aligning with principal stress directions reduce embodied carbon, and switching to timber is beneficial) while giving more precise quantifications of these relative decision impacts. The density variable leads to results that are potentially counterintuitive: in all cases, increasing the density (number of subdivisions) of the grid worsens performance. This result is likely due to the discrete nature of available cross sections. Our method can thus expand an experience-based understanding of the design problem based on data. This is especially useful in high-dimensional design spaces.

One can also understand the relative performance of the reference design with respect to its local neighborhood: for Design 4 in Fig. 7d, several design decisions have a large negative sensitivity, indicating that there are a variety of ways to significantly improve performance. We can more rigorously understand decision-making insights in the local neighborhood by evaluating the sensitivity of designs within a neighborhood of the original variable space of the reference design. An example is shown (on one of Fig. 7's reference designs) in Fig. 8. Relative magnitudes of sensitivity within each neighborhood are more meaningful than relative or absolute magnitudes of sensitivity across neighborhoods due to scaling of categorical sensitivities.

Examining neighborhood sensitivities contextualizes the degree to which local insights can be generalized in the larger neighborhood. For example, we can consider Design 1 across Fig. 7 and Fig. 8. The local sensitivities in Fig. 7a show that for Design 1, switching from current topology C to topologies B or F reduce embodied carbon. However, Fig. 8 indicates that in the specified neighborhood, this is almost always true for a switch to F but not for B.

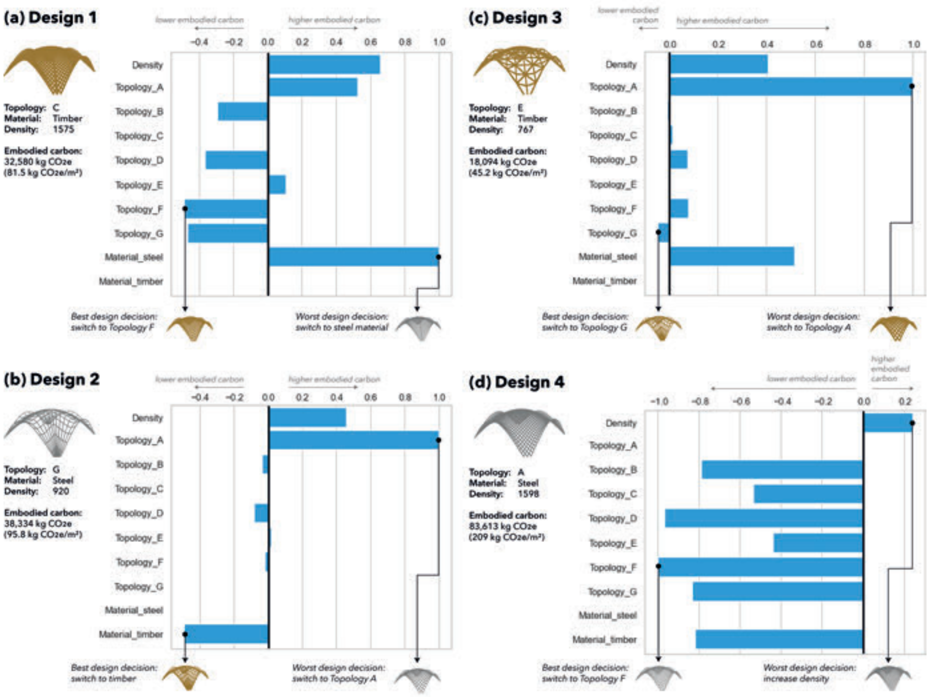

Fig. 7: Local sensitivities of 4 reference designs, annotated with design interpretation. The sensitivities of the categorical variables are scaled to match the order of magnitude of the continuous sensitivity individually for each design, so the magnitudes of sensitivity are not directly comparable across designs.

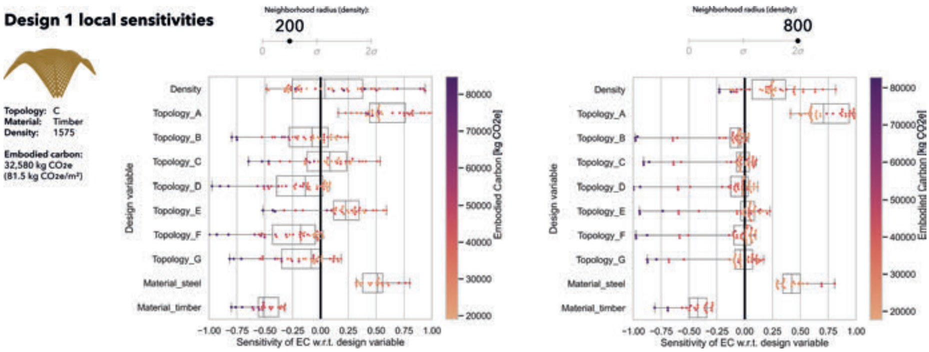

Fig. 8: Local sensitivities of a reference design, in neighborhoods of varying sizes. The radius is applied to the continuous variable density, while a full factorial sampling is done for the categorical variables. The number of designs does not increase with increasing neighborhood radius because the same number of designs are sampled (random uniform) regardless of neighborhood size.

In general, computing local sensitivities – either at a reference design (Fig. 7) or in the neighborhood around it (Fig. 8) – already improve on EDA by giving specific decision-making feedback.

4.2 Global sensitivities: Data-driven rules of thumb

Sensitivities at the global scale provide more generalized insights on decision-making for the design problem. These generalized insights are similar to “rules of thumb” typically employed by designers during more manual design. We extend the box-and-whisker swarm plots used by Balmer, Kuhn et al. 2022, enhancing the decision-making process with performance filtering. In Fig. 9, we illustrate this approach in the context of the gridshell case study. By adjusting the performance filter, we control the number of designs included in the analysis. Using the cVAE decoder, we generate additional designs that align with the filtered performance criteria, along with their corresponding sensitivities. Focusing on the top 5 % performing designs (Fig. 9d) allows for more targeted insights regarding influential design decisions among already high-performing designs, compared to considering the entire design space (Fig. 9a).

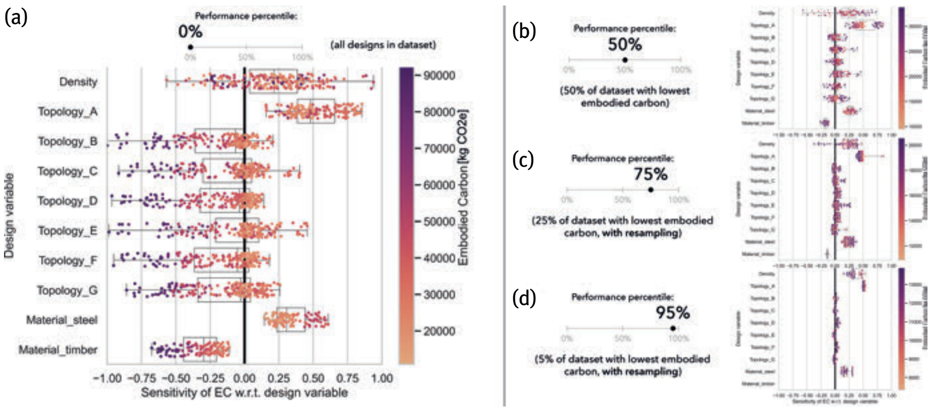

Fig. 9: Global sensitivities of the dataset at varying performance filters.

4.3 Quantifying the influence of design variables

The Gaussian Mixture Model (GMM) fitting approach described in Sec. 3.3 can be applied to any distribution of sensitivities; for brevity, we only show results on global sensitivities. The resulting means, standard deviations, and relative weights determined by the GMM fitting are shown in Figure 10. This representation reduces some of the complexity of the same distribution’s box-and-whisker swarm plot (Fig. 9a).

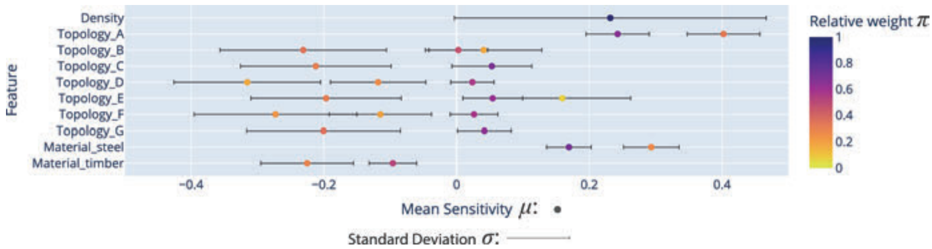

Fig. 10: Results of the GMM fitting on global sensitivity distributions determined by the cVAE. For many decisions, the data are distributed in multiple modes, which is captured in the GMM fit and represented here graphically. This indicates that the design decision can have a variety of impacts on performance, depending on location in the design space.

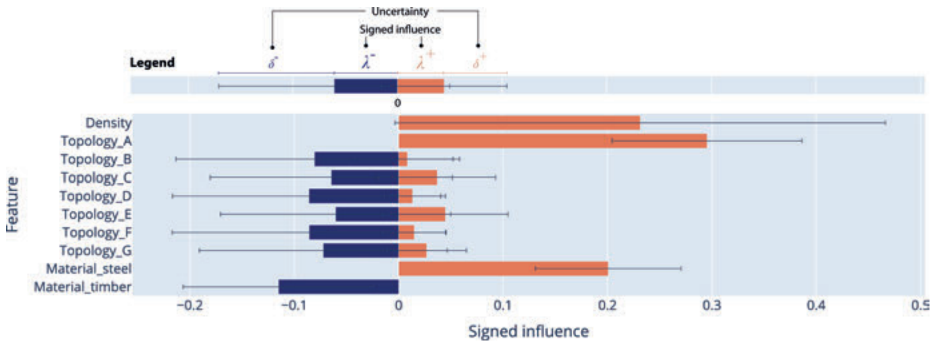

Fig. 11: Signed influence metrics of design variables on embodied carbon.

Signed influence metrics are provided in Fig. 11, computed by combining the results of the GMM model for each decision (Fig. 4). Negative influence metrics λ^- indicate a variable's influence on enhancing performance (lower embodied carbon), while positive influence metrics λ^+ on worsening performance (higher embodied carbon). δ^- and δ^+ represent the associated uncertainty metrics for each influence metric, respectively. By identifying variables with higher λ^+ and lower δ^+ values across the design space, one can understand that switching to Topology A overwhelmingly (in magnitude and certainty) worsens performance, followed closely by switching to steel. Switching to timber is the most influential design decision (in magnitude and likelihood) for improving performance, though not by much more than switching to topologies other than A. Increasing density generally worsens performance but with high uncertainty. We can confirm this by revisiting the EDA (Fig. 6) and noticing that changes in embodied carbon are relatively flat for a given topology of increasing density, an insight that was harder to identify without Fig. 11.

The influence metrics of material categories steel and timber are unexpectedly asymmetric though expected to be symmetric (Fig. 6). The authors identified the cVAE

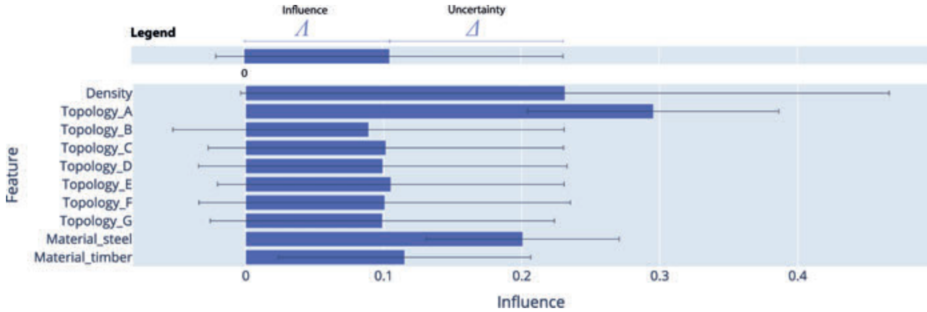

Fig. 12: Absolute influence metrics of design variables on embodied carbon.

Tab. 2: Table 2. Dashboard of absolute influence (categories aggregated into categorical variables).

	Density	Topology	Material
Aggregated absolute influence (Λ)	0.231	0.127	0.158
Aggregated uncertainty (Δ)	0.235	0.126	0.081

model encoder's prediction error to not be a zero-mean Gaussian, propagating into asymmetries in sensitivities and influence during the process of calculating λ. Solutions for overcoming this limitation are discussed in Sec. 5.

Finally, absolute influence metrics Λ and their associated uncertainty Δ are presented in Fig. 12. We can also present this as a "dashboard" (Tab. 2). These metrics further abstract previously available insights, approaching classic "rules of thumb", and might be useful for designers as general insights in the earliest stages of decision making. For example, the metrics indicate that choice of topology and density can be just as or more influential than material choice (Fig. 12), while material choice more certainly impacts performance than density or topology (Tab. 2).

4.4 Reflections

The example application of the proposed method demonstrated alignment with expected insights and rules of thumb. A summary of advantages and disadvantages associated with each visualization or metric is presented in Table 3. We observe tradeoffs between exhaustive representation of sensitivities and interpretability of sensitivities for decision-making insights. Overall, we recommend the GMM summary (Fig. 10) and signed influence metrics (Fig. 11) as representations that balance these properties.

Tab. 3: Parameters of gridshell case study. Output is an embodied carbon estimate under fully stressed design assumptions.

Visualization or metric	**Data represented**	**Advantages**	**Disadvantages**
Exploratory Data Analysis (Fig. 6)	Performance of full dataset	Understand range of performance available; Observe some prominent global trends	Unwieldy for multivariate design problems; Hard to get decision-making insights, especially locally
Bar plot (Fig. 7)	Sensitivities at 1 reference design	Get clear decision-making insights from the reference design	Insights are only applicable at a single reference design
Box-and-whisker swarm plot (Fig. 8, 9)	Distribution of sensitivities at multiple designs	Swarm representation exhaustively depicts the distribution	Can be hard to interpret swarm representation for decision-making; Box-and-whisker plot does not account for multi-modal behavior from local effects
Gaussian Mixture Model summary (Fig. 10)	Distribution of sensitivities at multiple designs	Simplifies complexity of swarm representation to a few key values; Accounts for possible multi-modal behavior from local effects	Reduced data resolution
Signed influence metrics (Fig. 11)	Distribution of sensitivities at multiple designs	Provides a quick snapshot on how each variable improves OR worsens performance across a set of designs	Drastically removes nuance from the GMM summary
Absolute influence metrics (Fig. 12)	Distribution of sensitivities at multiple designs	Most immediately interpretable; closest to a “rule of thumb”	No longer captures whether a decision improves or worsens performance

5 Conclusions

This paper proposes new metrics for quantifying the influence of design decisions, using sensitivities of a trained cVAE for architectural applications. The proposed method successfully serves as an automated, data-driven way to reveal insights for design decision-making. The cVAE is newly proposed for not only generating designs but also primarily extracting sensitivities of mixed variables in a design space, information that was not previously accessible from data generation alone.

The proposed influence metrics combine qualitative and quantitative representations to communicate actionable decision-making insights. We show a range of visualizations and metrics, discussing their respective limitations and advantages.

Ultimately, designers can choose levels of resolution and abstraction that best suit their design-decision-making process.

This approach has some limitations. The cVAE requires an upfront, time-consuming process of generating data and training (although it is instantaneous to query once trained). Model accuracy should be assessed in each application. Even models trained to reasonable accuracy provide approximate sensitivities; in our example, the model produced asymmetries in influence metrics that are expected to be symmetric. This can be addressed in future work by enforcing symmetric sensitivities or reducing prediction errors in the back-end model, and continuing to verify the intuition of the influence metric formulation in other design problems. Furthermore, the dataset in this case study contained a full factorial set of structurally viable designs; the cVAE is assumed to have learned to produce structurally viable designs. Future applications should consider an explicit representation or a condition for ensuring generated designs are structurally sound. Additionally, for the proposed influence metric, the GMM fitting necessarily assumes normal distributions for identifying multimodal behavior.

In future work, more elements of the cVAE can be harnessed for design, such as the latent space for correlation and neighborhood sampling, and the decoder for generating new design concepts. The method can be useful to inform and drive design problems beyond the case study presented here, such as investigating the material- and topology-related design decisions most influential in reducing carbon emissions in multi-story conventional buildings.

References

Balmer, V. M., S. V. Kuhn, R. Bischof, L. Salamanca, W. Kaufmann, F. Perez-Cruz, and M. A. Kraus. 2022. Design Space Exploration and Explanation via Conditional Variational Autoencoders in Meta-Model-Based Conceptual Design of Pedestrian Bridges. arXiv: 10.48550/arXiv.2211.16406

Bischof, R., and M. A. Kraus. 2022. Mixture-of-Experts-Ensemble Meta-Learning for Physics-Informed Neural Networks. In *Proceedings of 33. Forum Bauinformatik*.

Brown, N. C., and C. T. Mueller. 2018. Gradient-Based Guidance for Controlling Performance in Early Design Exploration. In *Proceedings of the International Association for Shell and Spatial Structures (IASS) Symposium*. Boston, MA.

Cho, S., and H.-S. Jung. 2003. Design Sensitivity Analysis and Topology Optimization of Displacement–Loaded Non-Linear Structures. *Computer Methods in Applied Mechanics and Engineering* 192 (22): 2539–53. DOI: 10.1016/S0045-7825(03)00274-3

Danhaive, R., and C. T. Mueller. 2021. Design Subspace Learning: Structural Design Space Exploration Using Performance-Conditioned Generative Modeling. *Autom. Constr.* 127: 103664. DOI: 10.1016/j.autcon.2021.103664

Fang, D., and C. Mueller. 2021. Flow-Informed Topology Design: Evaluating the Conformity of Structural Topologies with Vector Fields. In *Proceedings of the International Association for Shell and Spatial Structures (IASS) Symposium*. Surrey, UK.

Hu, Y., L. Anderson, T.-M. Li, Q. Sun, N. Carr, J. Ragan-Kelley, and F. Durand. 2020. DiffTaichi: Differentiable Programming for Physical Simulation. arXiv: 10.48550/arXiv.1910.00935

Jones, C., and Hammond, G. 2019. Inventory of Carbon & Energy v3.0. https://www.circularecology.com/embodied-energy-and-carbon-footprint-database.html

Regenwetter, L., A. Heyrani Nobari, and F. Ahmed. 2022. Deep Generative Models in Engineering Design: A Review. *J. Mech. Des.* 144 (7): 071704. DOI: 10.1115/1.4053859

Saltelli, A., M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola. 2008. *Global Sensitivity Analysis: The Primer.* John Wiley & Sons.