

Structural Interaction for Generative User Interfaces

Vincent Cavez
vcavez@stanford.edu
Stanford University
USA

Abstract

Generative user interfaces (genUI) promise to produce and adapt interface structures on-the-fly, yet current approaches lack a principled vocabulary for specifying how generated structures should behave during interaction. This paper argues that the central design challenge for genUI is not what structure to generate but how each structural rule should respond to user action. It presents Structural Interaction, a framework that defines structure as the set of rules in a directed graph of elements and rules, and characterizes rule behavior along two orthogonal dimensions: rigidity (how much rules can be shaped) and enforcement (how much rules can yield during interaction). Four values per dimension generate a 16-cell design space that enables designers and generative systems to reason precisely about when structure should constrain, adapt, or yield. We discuss how this framework reframes genUI design and outline the shifts HCI practice needs to support it.

CCS Concepts

• **Human-centered computing** → **Interaction design theory, concepts and paradigms**; *HCI design and evaluation methods*.

Keywords

structure, flexibility, malleable interfaces, interaction design, theoretical framework

ACM Reference Format:

Vincent Cavez. 2026. Structural Interaction for Generative User Interfaces. In *What does Generative UI mean for HCI Practice? Workshop of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 3 pages.

1 Motivation and Background

Productivity and creativity support tools rely on rules that govern how users create, organize, and transform content [11]. We call these rules *structure*. A spreadsheet’s grid, a score’s staff system, a design tool’s layer hierarchy are all structures. Structure simultaneously enables and constrains: the grid enables formulas and sorts precisely because it enforces a rigid schema, but that rigidity obstructs users when their intentions diverge from its logic [4]. This raises a tension between, on one hand, the power that structure provides through advanced manipulations, queries, and consistency checks; and, on the other hand, the freedom that users need to explore, iterate, and deviate from the system’s expectations. The same tension arises in music score writing, where notation programs

enforce structural rules that impede exploratory composition [7], calling for selective relaxation of those constraints [6].

As generative systems become capable of producing and adapting interface structures on behalf of users [3], the critical question shifts: how should generated structure behave during interaction? Several frameworks address this organizational layer. Assemblies in DIRA [1] describe what organizes representations and connects interaction techniques, but without specifying how those organizational elements respond when users push against them. Interaction substrates [9] define structured environments where users can adjust individual constraints through tweaking and create reusable configurations through templating; yet the question of how a constraint behaves when users push against it during interaction remains implicit. Malleable software [8, 10] demonstrates that users can reshape interface structure at runtime, but treats malleability as a platform-level property rather than a per-rule design parameter.

In the genUI space specifically, Cao et al. [3] propose generative interfaces driven by task-driven data models that an LLM generates and evolves, demonstrating that structure can drive interface generation. Cao et al. [2] design co-creation environments around compositional structures, addressing which structures to deploy and how to connect them. Their user evaluation confirms the value of making structure a primary design concern, but also reveals that when AI populates these structures, new tensions emerge: participants found the cost of structural revision rising as generated content accumulated, and high-fidelity AI output sometimes discouraged exploration and large-scale iteration. These findings point to a missing layer in the design vocabulary: there is no shared way to say “this rule should yield when the user pushes against it” or “this rule should be adjustable but not removable.”

2 Structural Interaction

The Structural Interaction framework [5] models the user interface as a directed graph containing two types of nodes: elements and rules. Elements are the entities users perceive: a spreadsheet cell, a note on a musical staff, a layer in a compositing tool. Elements are always targets in the graph, never sources: they do not trigger or constrain anything. Rules are the organizational nodes of the graph, and structure is defined as the set of all rules governing the organization of elements.

Rules are further divided into two types. Declarative rules are constraints that hold continuously and need no trigger: “cell B1 always displays the value of A1 multiplied by two,” “all elements in a layer inherit that layer’s opacity,” “this panel’s width never falls below 200 pixels.” Imperative rules are conditions that, if triggered, can perform an action: “if the user clicks this button, change its color,” “if a row is dragged past another, swap their positions.” This graph model distinguishes structure from content. Because elements are always targets, every transformation in the interface



This work is licensed under a Creative Commons Attribution 4.0 International License.
CHI '26, Barcelona, Spain
© 2026 Copyright held by the owner/author(s).

ultimately acts on or through rules: users never interact with elements directly; they interact with structure.

Within any structure, closely related rules cluster into sub-structures that users experience as a unit. A spreadsheet's grid logic is a sub-structure: a cluster of declarative and imperative rules governing how cells are created, addressed, referenced, and manipulated. Every rule carries a behavioral state defined by two orthogonal dimensions.

Rigidity captures how much rules can be shaped. Four values span a spectrum from system control to user control. A *fixed* rule is immutable: the email schema in Outlook requires a sender, a recipient, a subject, and a body, and no rule can modify this. A *negotiable* rule can be relaxed within bounds but not redefined: the volume limiter on iOS lets the user adjust the maximum level within a fixed range, but not redefine how limiting works. A *malleable* rule can be redefined entirely by the user, like a keyboard shortcut in Photoshop that can be reassigned to any function. An *authorable* rule goes further: the user can create rules that did not previously exist, as when a user writes a macro in Excel or defines custom geometric constraints in Fusion 360.

Enforcement captures how much rules can yield during interaction. Four values span a spectrum from unyielding to yielding. A *persistent* rule does not yield: a minimum panel width blocks hard at its limit. An *elastic* rule yields under pressure but restores itself: the magnetic timeline in Final Cut Pro resists gaps between clips, but clips can be forced apart before re-snapping on release. An *escapable* rule tolerates bypass that leaves a lasting result: indentation in VS Code enforces a tab size, but manual spacing can override it locally and the override persists. A *liftable* rule can be removed from enforcement entirely: spell-check in Word can be disabled for a document.

These two dimensions are orthogonal: any rigidity value can combine with any enforcement value. Desktop spreadsheet grids are (*fixed, escapable*): the grid logic itself is immutable, but paste-special operations and direct in-cell editing bypass dependencies and the results persist. The macOS Dock is (*negotiable, elastic*): size is adjustable within bounds, and magnification on hover yields then restores. The state couple thus generates a design space of 16 combinations.

3 Implications for Generative UI

The framework has three direct implications for genUI design.

GenUI must generate structural behavior, not just structural form. Current generative approaches focus on producing interface layouts and components [3]. The framework reveals that this is insufficient: for every rule a system generates, it must also specify that rule's rigidity and enforcement. A genUI system that deposits fully-formed content into a (*fixed, persistent*) structure produces an artifact users cannot modify or escape. This is consistent with the tensions Cao et al. [2] observed: as AI-generated content accumulated in their compositional structures, the cost of structural revision increased and high-fidelity output discouraged large-scale iteration. A system that generates (*negotiable, liftable*) structure would instead let users adjust generated constraints within bounds and selectively suspend them while exploring. The difference between a rigid, opaque genUI and a flexible, steerable one is not a

matter of "how much AI" but of which state couples the system assigns to its generated rules.

Consider a genUI system that generates a data dashboard. If the column layout rules are (*fixed, persistent*), the user cannot rearrange columns or break out of the generated grid. If they are (*malleable, escapable*), the user can redefine the layout entirely, and local overrides (manually resizing a column, dragging a widget outside the grid) persist as lasting deviations. The structural form is the same; the structural behavior is entirely different.

Structural states should align with workflow phases. Creative and analytical workflows move through phases with different structural needs. Exploration benefits from lower rigidity (*negotiable* or *malleable*) with *liftable* enforcement: the user can adjust or redefine constraints and suspend those that get in the way. Execution benefits from higher rigidity (*fixed*) with *persistent* enforcement: constraints hold firm so that the user can focus on content without structural interference.

A genUI system aware of workflow phase could dynamically shift rule states rather than imposing a single structural configuration throughout. During ideation, layout constraints might be (*negotiable, liftable*): adjustable and removable. As the user commits to a direction, the system could tighten to (*fixed, elastic*): the layout becomes fixed but still yields elastically under pressure, providing guidance without rigidity. During finalization, (*fixed, persistent*) locks the structure. This progression from loose to tight structure mirrors how designers and creative professionals naturally work, but current genUI systems have no vocabulary to express it.

The framework provides a shared specification language. For genUI to be steerable, designers and users need a vocabulary to communicate structural intent to the system. The state couple (rigidity, enforcement) provides a compact, domain-independent notation. A designer could specify that a generated layout's column constraints should be (*negotiable, elastic*) and its typography rules (*fixed, persistent*), giving the system precise behavioral targets rather than vague directives like "make it flexible." A user could request that generated constraints be "liftable" without knowing anything about the framework, and the system could interpret this as a shift along the enforcement dimension. The vocabulary scales from expert specification to natural-language interaction.

4 Discussion and Conclusion

Structural Interaction suggests that HCI and design practice will need to evolve in at least two ways to support genUI.

First, the unit of design must expand. Current practice centers on designing interactions with content: selecting, dragging, editing domain objects. When structure is generated rather than hardcoded, designers must also specify how each structural rule should respond to user action. This requires new design tools: notation systems for expressing structural states, simulation environments for previewing how a configuration behaves under user pressure, and evaluation methods that assess structural behavior independently of content quality.

Second, the relationship between user and structure must be renegotiated. In traditional interfaces, structure is the designer's province: users interact through structure, not with it. GenUI dissolves this boundary. If an LLM can generate and evolve structural

rules [3], users need mechanisms to inspect, override, and author those rules (shifting along the rigidity dimension from *fixed* toward *authorable*). The framework clarifies that “user control over AI-generated interfaces” is not a single slider but at least two independent parameters per rule. A user who can override a generated constraint (*escapable*) but not redefine it (*fixed*) has a fundamentally different relationship to the system than one who can redefine constraints freely (*malleable*) but cannot escape them during interaction (*persistent*).

These claims call for empirical validation. Other modalities or domains may reveal additional values or new dimensions. However, we believe the core argument is sound: as interfaces become generative, the organizational schemas through which users act become a primary design material. Structure is not merely something interfaces have; it is something interfaces do, and designing genUI through structural behavior opens new possibilities for interfaces that are simultaneously powerful, flexible, and human-centered.

References

- [1] Joanna Bergström and Kasper Hornbæk. 2025. DIRA: A model of the user interface. *Int. J. Hum.-Comput. Stud.* 193, C (Jan. 2025), 14 pages. doi:10.1016/j.ijhcs.2024.103381
- [2] Yining Cao, Yiyi Huang, Anh Truong, Hijung Valentina Shin, and Haijun Xia. 2025. Compositional Structures as Substrates for Human-AI Co-creation Environment: A Design Approach and A Case Study. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 188, 25 pages. doi:10.1145/3706598.3713401
- [3] Yining Cao, Peiling Jiang, and Haijun Xia. 2025. Generative and Malleable User Interfaces with Generative and Evolving Task-Driven Data Model. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 686, 20 pages. doi:10.1145/3706598.3713285
- [4] Vincent Cavez, Caroline Appert, and Emmanuel Pietriga. 2024. Spreadsheets on Interactive Surfaces: Breaking through the Grid with the Pen. *ACM Trans. Comput.-Hum. Interact.* 31, 2, Article 16 (Jan. 2024), 33 pages. doi:10.1145/3630097
- [5] Vincent Cavez, Kashif Imteyaz, and Anne-Flore Cabouat. 2026. Structural Interaction: Shifting the Focus of User Interface Design. (2026). Under review.
- [6] Vincent Cavez, Catherine Letondal, Caroline Appert, and Emmanuel Pietriga. 2025. EuterPen: Unleashing Creative Expression in Music Score Writing. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 760, 16 pages. doi:10.1145/3706598.3713488
- [7] Vincent Cavez, Catherine Letondal, Emmanuel Pietriga, and Caroline Appert. 2024. Challenges of Music Score Writing and the Potentials of Interactive Surfaces. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 728, 16 pages. doi:10.1145/3613904.3642079
- [8] Clemens N. Klokmoose, James R. Eagan, Siemen Baader, Wendy Mackay, and Michel Beaudouin-Lafon. 2015. Webstrates: Shareable Dynamic Media. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (Charlotte, NC, USA) (UIST '15)*. Association for Computing Machinery, New York, NY, USA, 280–290. doi:10.1145/2807442.2807446
- [9] Wendy E. Mackay and Michel Beaudouin-Lafon. 2025. Interaction Substrates: Combining Power and Simplicity in Interactive Systems. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 687, 16 pages. doi:10.1145/3706598.3714006
- [10] Philip Tchernavskij. 2019. *Designing and Programming Malleable Software*. Ph.D. Dissertation. <http://www.theses.fr/2019SACL499>
- [11] Jenifer Tidwell. 2010. *Designing Interfaces*. O'Reilly Media, Inc.