

A Design Space for Peripheral Interaction: Evidence Mapping and Transferable Implications

Haopeng Wu and Wonseok Yang

Shibaura Institute of Technology, 3-7-5 Toyosu, Koto-ku, Tokyo, Japan

ABSTRACT

Peripheral interaction research spans ambient displays, glanceable mobile cues, and emerging spatial/XR mechanisms, but design knowledge remains fragmented across devices and contexts. We present a two-stage synthesis that constructs a unified design space and derives cross-device transferable implications. First, we perform macro-level evidence mapping using two primary axes—attention demand (D1) and spatial placement (D10)—and locate 188 strongly relevant studies (Tier 1) in a D1×D10 grid. The map reveals a highly skewed landscape dominated by background environmental systems, while sparse regions indicate underexplored opportunities for far-field and spatially anchored peripheral mechanisms. Second, we extract actionable transfer rules from a deep set of 43 papers selected through maximum-difference coverage sampling across the D1×D10 grid. We apply a 12-dimension codebook (D1–D12) covering signal, interaction, temporal, and governance properties. Cross-dimensional coupling analysis reveals recurring tension structures (e.g., detectability vs. disruptiveness, aesthetics vs. legibility, and adaptivity vs. user control) and supports seven reusable implications expressed as operational transfer rules. Together, the unified design space, evidence map, and transfer rules provide a systematic foundation for cross-device peripheral interaction design and future periphery-aware HCI research.

Keywords: Peripheral interaction, Design space, Evidence map, Cross-device UX

INTRODUCTION

Peripheral interaction aims to support user awareness and action without continuously pulling attention away from primary activities (Weiser, 1991; Weiser and Brown, 1997). Across HCI, related work has emerged under multiple labels—including calm technology, ambient information systems, glanceable interfaces, and more recent spatial/XR notification concepts—often with device- and modality-specific assumptions (Heiner et al., 1999; Pousman and Stasko, 2006; Bakker and Niemantsverdriet, 2016; Lu, 2021). As interactive ecosystems become increasingly multi-device, designers face a persistent gap: how to transfer peripheral interaction principles from one device class (e.g., environmental ambient artifacts) to another (e.g., wearables or spatially anchored XR) while preserving attentional balance, interpretability, and user control (Angelini et al., 2018; Lee et al., 2023). Prior research offers valuable concepts and case studies, but the knowledge

is dispersed. Two practical problems follow. First, it is difficult to understand where research density lies and where gaps remain, especially across attention demands and spatial configurations. Second, even when patterns are known within a device class, operational transfer rules across devices are rarely articulated in a reusable form. This paper addresses these problems via a two-stage synthesis with three contributions: 1. Unified design space: a 12-dimension design space for peripheral interaction, organized by primary axes D1 (attention demand) and D10 (spatial placement) and complemented by ten dimensions that enable cross-device transfer. 2. Evidence map: a D1×D10 evidence map of Tier 1 papers ($n = 188$) that characterizes density and gaps. 3. Reusable implications: seven cross-device implications expressed as operational transfer rules derived from deep full-text coding ($n = 43$). Figure 1 previews our unified 12-dimension design space organized by D1×D10 and complemented by ten transfer dimensions.

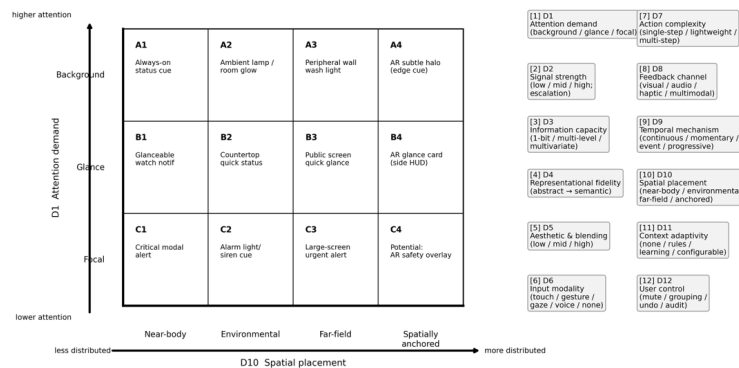


Figure 1: Unified 12D design space.

METHOD

Overview: Two-stage Synthesis

We conducted a two-stage synthesis to build a unified design space and derive transferable implications. Stage 1 performs macro-level evidence mapping using lightweight coding on D1 (attention demand) and D10 (spatial placement), grounded in prior peripheral interaction and attention-continuum discussions (Bakker and Niemantsverdriet, 2016; Angelini et al., 2018). Stage 2 performs deep full-text coding on a purposefully selected subset to identify cross-dimensional couplings, tension structures, and reusable implications across device classes and interaction contexts.

Search Strategy and Databases

We searched three scholarly databases (Web of Science Core Collection, Scopus, IEEE Xplore) using a convergent query set covering peripheral interaction and adjacent concept clusters (e.g., peripheral display, ambient information systems, calm technology, glanceability, interruption

management, and XR notifications). Searches were executed and exported in January 2026. Records were exported in RIS/CSV formats and imported into Rayyan for screening and de-duplication.

Screening, De-duplication, and Tiering

Across the three databases, we identified 2,398 records. After de-duplication in Rayyan, 1,577 unique records remained. We conducted title/abstract screening to retain works relevant to peripheral interaction mechanisms (broadly construed as interactions intended to operate under background/glance attention with limited disruption), yielding 700 records for tiering. Lightweight coding then assigned these records into Tier 1 ($n = 188$) for strong relevance (used for evidence mapping) and Tier 2 ($n = 81$) for moderate relevance (retained as a backup pool for deep-set selection). The remaining records ($n = 431$) were excluded after tiering because they did not provide sufficient fit to the operational scope (e.g., peripheral concepts mentioned without explicit interaction mechanisms, unclear attentional framing, or insufficient design/interaction detail for mapping).

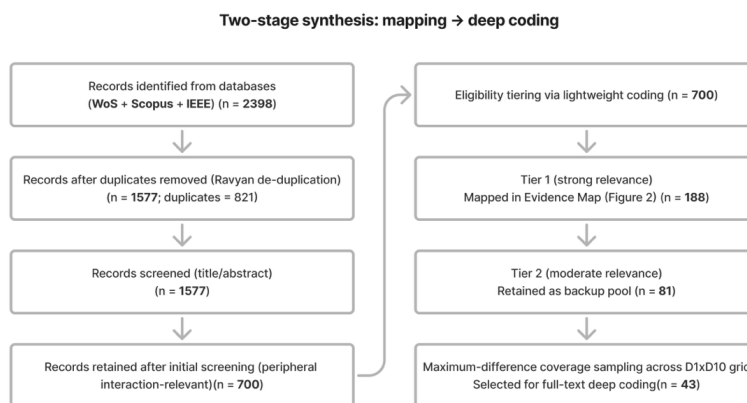


Figure 2: PRISMA-style flow.

Evidence Mapping (D1×D10)

We mapped Tier 1 papers onto the D1×D10 grid. D1 includes Background / Glance / Focal, and D10 includes Near-body / Environmental / Far-field / Spatially anchored. The map (Figure 3) summarizes density and gaps across attention and spatial placement.

Deep Set Selection and Full-Text Coding (n = 43)

To derive cross-device transfer rules, we selected a deep set of papers via maximum-difference coverage sampling across the D1×D10 grid, aiming to cover dense regions (dominant patterns) and sparse regions (boundary cases). We retrieved and compiled full-text PDFs, yielding a final deep set of 43 papers for full coding using a 12-dimension codebook (D1–D12) spanning signal properties, interaction modality, temporal mechanisms, spatial placement,

adaptivity, and user control. Cross-dimensional couplings were synthesized by identifying recurring co-occurrence patterns and abstracting them into operational implications and tension structures, which are reported in the Results section and summarized in Table 1.

RESULTS

Deep Set Characteristics (n = 43)

To derive cross-device transferable implications beyond macro-level mapping, we constructed a deep set of 43 papers for full-text coding. The deep set spans 1999–2024, with most papers published in the 2010s (28/43, 65%), followed by the 2020s (8/43, 19%). This pattern suggests that peripheral interaction research matured primarily in the last decade. Based on DOI-derived publisher inference, the corpus is dominated by HCI publication channels. ACM accounts for the largest share (28/43, 65%), followed by Springer (6/43, 14%) and IEEE (3/43, 7%), with the remainder distributed across other outlets. Where venue names are available, CHI Proceedings is the most frequent venue (n = 10), followed by ACM International Conference Proceeding Series (n = 4), with smaller contributions from journals such as International Journal of Design (n = 2) and Personal and Ubiquitous Computing (n = 2). Using the coarse device-domain labels in our coding, the deep set covers ambient systems (23/43, 53%), mobile/wearable contexts (12/43, 28%), XR/spatial computing (5/43, 12%), and desktop contexts (3/43, 7%). Evidence types include lab studies (15/43, 35%), field deployments (12/43, 28%), and conceptual/framework contributions (16/43, 37% combined), supporting both empirical findings and design-knowledge synthesis.

D1×D10 Evidence Map: Density and Gaps

Figure 3 maps the Tier 1 corpus (n = 188) onto attention demand (D1) and spatial placement (D10), revealing a highly skewed distribution. The dominant cluster lies in Background × Environmental (n = 114; ~61%), indicating that peripheral interaction research remains strongly anchored in environment-embedded, low-demand ambient systems. A secondary mass appears in Near-body × (Background/Glance) (n = 35; ~19%), suggesting growing adoption of peripheral mechanisms on mobile and wearable devices, largely for lightweight awareness and glanceable updates. In contrast, multiple regions remain underpopulated: Focal × Spatially Anchored has no records (n = 0), Glance × Environmental is nearly absent (n = 1), and Background × Far-field is sparsely explored (n = 4). These gaps point to opportunities for extending peripheral principles to spatial computing and large-display contexts while preserving attentional balance. This Tier 1 map provides the macro-level landscape. To derive cross-device transferable implications, we subsequently selected a deep set of 43 papers via maximum-difference coverage sampling across the D1×D10 grid for full-text coding and cross-dimensional coupling analysis.

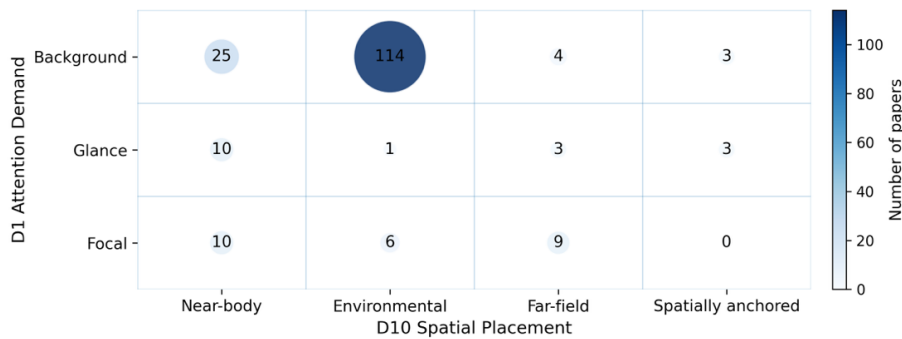


Figure 3: D1xD10 evidence map, Tier 1 n = 188.

Cross-Dimensional Couplings and Tension Structures (Deep Set, n = 43)

Cross-dimensional coupling analysis of the deep set (n = 43) indicates that D1 (attention demand) and D10 (spatial placement) function as organizing axes, but they do not operate independently. Instead, they co-vary with representational encoding, temporal behavior, interaction cost, and control mechanisms. This pattern aligns with foundational discussions of calm technology and ambient information systems, while also matching later attention-continuum perspectives that treat peripheral interaction as a dynamic shift rather than a fixed interaction mode (Weiser and Brown, 1997; Pousman and Stasko, 2006; Bakker and Niemantsverdriet, 2016). To make the couplings reusable across device classes, we summarize the seven implications as compact transfer rules in Table 1, including the coupled dimensions, expected transfer targets, and minimal evaluation focus for each rule.

Attention–Space Alignment (D1xD10 as Transfer Anchor)

Across devices, spatial placement strongly conditions what level of attention demand is sustainable. Environmental embedding supports persistent background awareness, whereas far-field and spatially anchored cues more easily drift toward focal capture unless managed by low-information signaling and progressive reveal. This yields the primary transfer rule: fix a target D1xD10 region first, then tune signal and timing (D2/D9) to stay within the intended attentional budget (Table 1, Implication 1 & 6).

Signal–Time Shaping (D2/D3/D4 ↔ D9)

Always-on peripheral awareness tends to pair with low information capacity and more abstract representation to avoid overload and habituation. When higher urgency is required, systems commonly employ escalation ladders: begin with subtle cues and increase intensity or modality only when conditions persist or risk rises. This tension—detectability vs. disruptiveness—recurs across mobile notifications, ambient artifacts, and XR cues (Table 1, Implication 2 & 3).

Table 1: Reusable implications: Deep set, n = 43. Abbreviations: BG = Background, GL = Glance, FO = Focal; NB = Near-body, ENV = Environmental, FF = Far-field, SA = Spatially anchored.

Implication	Coupling (Dims)	Operational Transfer Rule (Compact)	Where it Transfers (Tags)	Evaluation Focus (Minimal Set)
I1 Anchor first	D1×D10 (+D9)	Porting across devices starts by fixing target D1×D10 cell; then tune timing (D9) to keep the same “attention signature” (BG/GL/FO).	NB/ENV/FF/SA; cross-device handoff	Primary-task cost; awareness accuracy; perceived disruption
I2 Persistent = abstract	D1(BG)+D9 (continuous)+D4 (abstract) (+D3 low)	For always-on awareness, keep info capacity low (1-bit/low-level) and representation abstract; reveal semantics only on demand.	ENV/NB; ambient & wearables	Long-term interpretability; habituation; misinterpretation rate
I3 Escalation ladder	D2+D8+D9 → D1 shift	Use progressive cueing: start subtle, amplify only when conditions persist/urgency rises (intensity/modality/time). Avoid high-salience first.	Notifications; safety cues; XR alerts	Detectability vs disruptiveness; response time; misses/false alarms
I4 Adaptivity needs governance	D11×D12 (±D8 social exposure)	More intelligence/automation requires stronger user control: mute/schedule/granularity/override + inspectability/auditability.	MOB/Wearable/Smart-home; shared spaces	Trust; perceived autonomy; privacy comfort; preference stability
I5 Keep periphery = low interaction cost	D6×D7 (+D1)	To remain peripheral, reduce input/action complexity before increasing content richness: no-input or single-step > multi-step drill-down.	Hands-busy contexts; XR; driving/field work	Interference; error types; learnability under workload

(Continued)

Table 1: Continued.

Implication	Coupling (Dims)	Operational Transfer Rule (Compact)	Where it Transfers (Tags)	Evaluation Focus (Minimal Set)
I6 Far-field tends to focal—constrain it	D10(FF) risk → D1(FO); mitigated by D2+D3+D9	For far-field/large displays, keep capacity low, place cues in peripheral regions, and use gradual onset; reserve details for user-initiated access.	Public displays; meeting rooms; large screens	Noticeability vs distraction; spatial legibility; group disruption
I7 Aesthetics vs legibility	D5×D4 (+D12)	When blending cues aesthetically/unremarkably, add onboarding/examples + consistent mappings + optional semantic reveal to avoid ambiguity/uncanny effects.	Decorative ambient artifacts; home/office acceptance	Aesthetic acceptance; mapping learnability; long-term engagement

Interaction Workload Under Ongoing Primary Tasks (D6/D7 With D1)

Peripheral designs that aim to remain in background or glance attention often minimize input and action complexity. In hands-busy or safety-critical contexts, “no-input” or single-step acknowledgement patterns appear repeatedly, while multi-step interaction sequences are strongly associated with focal attention capture. This coupling supports a pragmatic transfer rule: to keep periphery, reduce interaction cost before increasing intelligence or content richness (Table 1, Implication 5).

Governance, Acceptability, and Trust (D11/D12; With D5/D4)

As systems become adaptive (e.g., intelligent notification management) or more socially visible (e.g., ambient displays in shared spaces), user control becomes a prerequisite for acceptability. Stronger modalities and automated decision policies require stronger governance features—mute, grouping, undo, configuration, and auditability. A related tension arises between aesthetic blending and semantic legibility; providing optional legibility layers and control mitigates misinterpretation and privacy concerns (Table 1, Implication 4 & 7).

Together, these couplings operationalize cross-device transfer as “anchor on D1×D10, then tune signal, time, interaction cost, and governance to preserve attentional balance.”

DISCUSSION

What the Evidence Map Implies for Cross-Device Research

The evidence map shows that peripheral interaction research remains disproportionately concentrated in low-attention, environmental placements (background/ambient room-scale displays), while fewer studies address mid-attention and body-/device-proximal configurations. This skew is consistent with the field's historical roots in calm technology and ambient information systems, and with later efforts to formalize attention as a continuum rather than a binary state (Weiser and Brown, 1997; Pousman and Stasko, 2006; Bakker and Niemantsverdriet, 2016). More recent work expands the map toward situated and body-proximal peripheral mechanisms, but coverage remains uneven across spatial placements and transitions between attention levels (An et al., 2019; Vatavu, 2017; Lu, 2021).

Toward Reusable Transfer Rules Rather Than Device-specific Heuristics

The deep-set synthesis suggests that transferable implications should be expressed as cross-dimensional rules rather than device-specific heuristics. Across the sampled literature, successful designs repeatedly manage the same tension structures while instantiating them with different media and placements.

- 1) Balance detectability with disruptiveness. Notification salience should be calibrated by combining cue modality, intensity, and onset timing rather than by increasing intensity alone. Work on gaze-contingent notifications and response-time tradeoffs shows that “more noticeable” cues can quickly become performance-degrading if timing and placement are poorly tuned (Klauck et al., 2017; Löcken et al., 2017). This also appears in recent ambient-to-disruptive notification studies, where gradual escalation and interpretable transitions help users maintain trust and control (Lee et al., 2023).
- 2) Separate aesthetic embedding from semantic precision. Ambient and peripheral interfaces benefit from aesthetic integration because it supports long-term acceptance and background presence, but semantic decoding still requires legibility when users shift attention toward the display. Classic ambient systems and later taxonomies already point to this tradeoff, and recent social or lifestyle-oriented ambient feedback work suggests that decorative or metaphorical forms work best when paired with interpretable signal mappings (Heiner et al., 1999; Pousman and Stasko, 2006; Choi et al., 2016).
- 3) Support progressive disclosure and user-adjustable escalation. Peripheral systems should not force a single fixed notification level. Instead, they should support transitions from ambient awareness to focused inspection, with optional controls for sensitivity or escalation. This pattern is visible in glanceable feedback systems and in safety-related notification contexts, where user trust depends on being able to interpret and regulate how alerts surface (Gouveia et al., 2016; Papachristos et al., 2020; Lee et al., 2023).

- 4) Preserve spatial congruence and bodily relevance. Cross-device transfer works better when placement and motion cues remain meaningful relative to the user's body, task space, or environmental locus. Studies on spatially distributed ambient information, wearable ambient displays, and peripheral attention-shifting mechanisms all suggest that spatial congruence reduces interpretation cost and improves response reliability in divided-attention conditions (An et al., 2019; Li et al., 2020; Gruenefeld et al., 2018).
- 5) Match temporal behavior to task rhythm. Peripheral cues should be designed as temporal phenomena (e.g., persistence, fade-in, repetition, decay), not just static visual forms. Evidence from dual-task and attentional studies indicates that onset profile and cue duration strongly influence both noticeability and interference, especially in time-sensitive contexts (Stewart and Billingham, 2016; Klauck et al., 2017; Löcken et al., 2017).
- 6) Couple adaptivity with transparent control. Adaptive peripheral behavior can reduce burden, but opaque adaptation risks confusion or mistrust. The literature suggests that adaptation is most acceptable when users can inspect, override, or personalize the system's behavior, particularly in always-on or context-sensitive environments (Angelini et al., 2018; Gouveia et al., 2016; Papachristos et al., 2020).
- 7) Translate interaction principles, not surface forms. The most reusable knowledge across devices lies in transferable relations (e.g., salience-escalation, spatial congruence, temporal fit, user control), not in copying specific visual or physical forms. This is especially important when moving from ambient artifacts to wearables or XR, where display technology changes but attention-management problems remain structurally similar (Bakker and Niemantsverdriet, 2016; Lu, 2021; Li et al., 2020).

Evaluation Guidance

Across implications, recurring evaluation foci include primary-task interference (time, errors), awareness accuracy, detectability–disruption trade-offs, and perceived autonomy/trust. We recommend aligning evaluation measures with the targeted D1 level: background designs prioritize low intrusion and long-term interpretability, while focal designs require task performance and safety-oriented metrics.

LIMITATIONS AND FUTURE WORK

This synthesis is constrained by database coverage and by how the coding dimensions were operationalized. Although the two-stage approach balances landscape mapping (Tier 1) and deep inference ($n = 43$), the deep set still depends on full-text availability and may underrepresent inaccessible, non-indexed, or non-English work. In addition, the D1×D10 evidence map relies on title/abstract-based lightweight coding, which may blur borderline cases.

Future work should address these limits in three ways. First, the corpus can be expanded to additional databases and gray literature where relevant. Second, coding reliability can be strengthened through multi-coder agreement

and adjudication protocols. Third, the unified design space can be extended to cover multi-device orchestration patterns (e.g., handoff, redundancy, and cross-surface semantics) as spatial computing ecosystems mature.

CONCLUSION

We presented a unified 12-dimension design space for peripheral interaction and a two-stage synthesis that bridges evidence mapping with operational transfer rules. A D1×D10 evidence map of Tier 1 papers (n = 188) reveals a landscape dominated by background environmental systems and highlights underexplored regions for far-field and spatially anchored designs. Deep full-text coding of 43 papers identifies cross-dimensional couplings and tension structures, summarized as seven reusable implications. Together, the design space, evidence map, and transfer rules provide a systematic foundation for cross-device peripheral interaction design and future periphery-aware HCI research.

REFERENCES

- An, P., Bakker, S., Ordanovski, S., Taconis, R., Paffen, C.L.E., and Eggen, B. (2019). Unobtrusively enhancing reflection-in-action of teachers through spatially distributed ambient information. *Conference on Human Factors in Computing Systems - Proceedings*. doi:10.1145/3290605.3300321.
- Angelini, L., Mugellini, E., Abou Khaled, O.A., Couture, N., van den Hoven, E., and Bakker, S. (2018). Internet of tangibles: Exploring the interaction-attention continuum. doi:10.1145/3173225.3173334.
- Bakker, S. and Niemantsverdriet, K. (2016). The interaction-attention continuum: Considering various levels of human attention in interaction design. *International Journal of Design*, 10(2), 1–14.
- Choi, S., Jeong, H., Lee, U., and Ko, M. (2016). LockDoll: Providing ambient feedback of smartphone usage within social interaction. *Conference on Human Factors in Computing Systems - Proceedings*, 7.0, 1165–1172. doi:10.1145/2851581.2892445.
- Gouveia, R., Pereira, F., Karapanos, E., Munson, S.A., and Hassenzahl, M. (2016). Exploring the design space of glanceable feedback for physical activity trackers. doi:10.1145/2971648.2971754.
- Gruenefeld, U., Stratmann, T.C., Jung, J., Lee, H., Choi, J., Nanda, A., and Heuten, W. (2018). Guiding Smombies: Augmenting Peripheral Vision with Low-Cost Glasses to Shift the Attention of Smartphone Users. doi:10.1109/ISMAR-Adjunct.2018.00050.
- Heiner, J.M., Hudson, S.E., and Tanaka, K. (1999). The information percolator: Ambient information display in a decorative object. *Conference on Human Factors in Computing Systems - Proceedings*, 1.0, 141–148. doi:10.1145/320719.322595.
- Klauck, M., Sugano, Y., and Bulling, A. (2017). Noticeable or distracting? A design space for gaze-contingent user interface notifications. *Conference on Human Factors in Computing Systems - Proceedings*, 1779–1786. doi:10.1145/3027063.3053085.
- Lee, J.G.W., Lee, B., and Choe, E.K. (2023). Decorative, Evocative, and Uncanny: Reactions on Ambient-to-Disruptive Health Notifications via Plant-Mimicking Shape-Changing Interfaces. *Conference on Human Factors in Computing Systems - Proceedings*. doi:10.1145/3544548.3581486.

- Li, H., Jarusriboonchai, P., Muller, H., Harjuniemi, E., and Häkkinen, J. (2020). Emotional Communication between Remote Couples: Exploring the Design of Wearable Ambient Displays. *ACM International Conference Proceeding Series*. doi:10.1145/3419249.3420139.
- Lu, F. (2021). [DC] Glanceable AR: Towards an always-on augmented reality future. doi:10.1109/VRW52623.2021.00241.
- Löcken, A., Blum, S., Stratmann, T.C., Gruenefeld, U., Heuten, W., Boll, S., and van de Par, S. (2017). Effects of location and fade-in time of (audio-)visual cues on response times and success-rates in a dual-task experiment. doi:10.1145/3119881.3119894.
- Papachristos, E., Merritt, T.R., Jacobsen, T., and Bagger, M.P. (2020). Hear me out: The role of voice in adding human-like perceptual qualities to a familiar object. *ACM International Conference Proceeding Series*, 59–70. doi:10.1145/3428361.3428400.
- Pousman, Z. and Stasko, J. (2006). A taxonomy of ambient information systems: Four patterns of design. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '06)*, 67–74. doi:10.1145/1133265.1133277.
- Stewart, J. and Billingham, M. (2016). A wearable navigation display can improve attentiveness to the surgical field. *International Journal of Computer Assisted Radiology and Surgery*, 11.0(6.0), 1193–1200. doi:10.1007/s11548-016-1372-9.
- Vatavu, R.-D. (2017). Smart-Pockets: Body-deictic gestures for fast access to personal data during ambient interactions. *International Journal of Human Computer Studies*, 103.0, 1–21. doi:10.1016/j.ijhcs.2017.01.005.
- Weiser, M. (1991). The computer for the 21st century. *Scientific American*, 265(3), 94–104.
- Weiser, M. and Brown, J. S. (1997). *The coming age of calm technology*. Springer.