

Abstract

We investigate whether one semantic scholarly document can produce coherent renditions for print, e-ink, and continuous reading while preserving identity, relationships, reading position, and annotations.

[sec-proposition] 1. Proposition

[p-proposition-1] Documents are usually flattened into pages too early. Once meaning becomes coordinates, every new screen or sheet becomes a repair job. Semantic Responsive Typesetting reverses that relationship: the document stays persistent; pages, columns, and line breaks are computed for a target.

[q-core] One semantic document. Multiple computed renditions. Persistent identity, relationships, reading position, and annotations.

[sec-model] 2. Composition model

[p-model-1] A canonical graph is combined with a target profile, a deterministic composition policy, and optional target-specific overrides. The renderer selects component variants, measures flow, records violations, and emits both the rendition and an inspectable manifest.

Print presentation: full-span-emphasized

[fig-pipeline] Canonical composition pipeline

[cap-pipeline] Stable semantic content flows through target constraints; geometry is output, never source truth.

[sec-gap] 3. Apparent research gap

[p-gap-1] Adaptive layout, PDF reflow, scholarly reconstruction, and robust annotation anchoring each have substantial prior art. The credible contribution is their integration into a fidelity-constrained, finite-height system and its evaluation across materially different reading surfaces.

[sec-method] 4. First experiment

[p-method-1] The first experiment holds content constant and varies representation geometry: A4 print, reMarkable Paper Pro, reMarkable Paper Pro Move, and continuous mobile. We test structural preservation, layout quality, anchor continuity, and composition time.