



HOCHSCHULE OSNABRÜCK
UNIVERSITY OF APPLIED SCIENCES

COMPARATIVE STUDY OF BLOCKCHAIN VISUALIZATION STRATEGIES ON MOBILE DEVICES

Fakultät für Ingenieurwissenschaften und Informatik
der Hochschule Osnabrück
Studiengang “Informatik – Medieninformatik”

Bachelorarbeit

zur Erlangung des akademischen Grades
Bachelor of Science

vorgelegt von

Finn Marten Garrels

Matrikelnummer:	1051281
Ausgabedatum:	09.01.2026
Abgabedatum:	05.03.2026
Erstprüfer:	Prof. Dr. Heinz-Josef Eikerling
Zweitprüfer:	Roman Dovgopol

Acknowledgements

I would like to thank the entire FiDi Tech team for their continuous support throughout the project phase. Their commitment to engage with early design iterations, provide honest feedback, and participate in the evaluation made this thesis possible. I am equally grateful to all additional study participants who volunteered their time and attention for the comparative evaluation.

Abstract

Blockchain ecosystems produce temporal, categorical, and hierarchical performance data that must increasingly be accessed on mobile devices. Despite this demand, no empirical study has compared alternative visualization strategies for blockchain network data on smartphones, and the intersection of mobile visualization constraints with the complexity of blockchain data remains unaddressed in published research.

This thesis designs, implements, and evaluates three structurally distinct visualization strategies for the same mobile blockchain data domain. Strategy A organizes data behind nested cards using progressive disclosure. Strategy B pairs charts with narrative text and search-based access. Strategy C distributes multiple chart types across swipeable panels at high information density. One strategy was realized as a native iOS application, two as interactive Figma prototypes. Ten participants evaluated all three in a within-subjects comparative study, using task scenarios, a modified System Usability Scale, experience ratings, and cross-strategy comparative instruments that measure usability and information clarity as separate constructs.

Progressive disclosure produced the strongest results across all instruments (SUS 78.0, ease of use 4.80, clarity 4.50). The high-density layout scored lowest on usability (SUS 59.5) while retaining comparatively more of its clarity rating (3.30 vs. ease of use 3.00). The most informative finding is a reversal in the relationship between the two dimensions: Strategy A is perceived as easier to use than it is clear, while Strategy C is perceived as clearer than it is easy to use. Usability and information clarity do not collapse onto a single quality axis but define a tradeoff space moderated by domain expertise.

The fidelity asymmetry between the native application and the Figma prototypes constitutes a methodological limitation: an unknown portion of Strategy A's advantage reflects implementation quality rather than visualization design. Within these boundary conditions, the thesis contributes a comparative evaluation framework that separates usability from information clarity, empirical evidence on mobile visualization tradeoffs in a real industry context, and documentation of the density gradient as a design variable whose optimal position depends on the user's expertise and analytical intent.

Zusammenfassung

Blockchain-Ökosysteme erzeugen zeitliche, kategoriale und hierarchische Leistungsdaten, die zunehmend auf mobilen Endgeräten abgerufen werden müssen. Trotz dieser Anforderung existiert keine empirische Studie, die alternative Visualisierungsstrategien für Blockchain-Netzwerkdaten auf Smartphones vergleicht. Die Schnittmenge aus mobilen Visualisierungseinschränkungen und der Datenkomplexität von Blockchain-Ökosystemen bleibt in der publizierten Forschung unbearbeitet.

Diese Arbeit entwirft, implementiert und evaluiert drei strukturell unterschiedliche Visualisierungsstrategien für dieselbe mobile Blockchain-Datendomäne. Strategie A organisiert Daten hinter verschachtelten Karten mittels Progressive Disclosure. Strategie B kombiniert Diagramme mit erläuterndem Text und suchbasiertem Zugang. Strategie C verteilt mehrere Diagrammtypen über wischbare Panels bei hoher Informationsdichte. Eine Strategie wurde als native iOS-Anwendung umgesetzt, zwei als interaktive Figma-Prototypen. Zehn Teilnehmer bewerteten alle drei Strategien in einem Within-Subjects-Vergleich anhand von Aufgabenszenarien, einer modifizierten System Usability Scale, Erfahrungsbewertungen und strategieübergreifenden Vergleichsinstrumenten, die Usability und Informationsklarheit als getrennte Konstrukte erfassen.

Progressive Disclosure erzielte die stärksten Ergebnisse über alle Instrumente hinweg (SUS 78,0; Bedienbarkeit 4,80; Klarheit 4,50). Das Layout mit hoher Informationsdichte schnitt bei der Usability am schwächsten ab (SUS 59,5), behielt jedoch einen vergleichsweise höheren Klarheitswert (3,30 gegenüber Bedienbarkeit 3,00). Der zentrale Befund ist eine Umkehrung im Verhältnis beider Dimensionen: Strategie A wird als leichter bedienbar wahrgenommen als klar, während Strategie C als klarer wahrgenommen wird als leicht bedienbar. Usability und Informationsklarheit bilden keine einzelne Qualitätsachse, sondern einen Tradeoff-Raum, der durch die Fachkenntnis der Nutzenden moderiert wird.

Die Fidelitätsasymmetrie zwischen der nativen Anwendung und den Figma-Prototypen stellt eine methodische Einschränkung dar: Ein unbekannter Anteil des Vorteils von Strategie A spiegelt Implementierungsqualität statt Visualisierungsdesign wider. Innerhalb dieser Rahmenbedingungen leistet die Arbeit drei Beiträge: ein vergleichendes Evaluationsframework, das Usability von Informationsklarheit trennt, empirische Evidenz zu mobilen Visualisierungs-Tradeoffs in einem realen Industriekontext sowie die Dokumentation des Dichtgradienten als Designvariable, deren optimale Position von der Expertise und der analytischen Absicht der nutzenden Person abhängt.

Contents

List of Figures	VI
List of Tables	VIII
List of Abbreviations	XI
1 Introduction	1
1.1 Problem Context and Motivation	1
1.2 Research Question and Objectives	2
1.3 Methodological Overview	3
1.4 Thesis Structure	4
2 Foundations	5
2.1 Blockchain Network Data	5
2.1.1 Network Architecture and Data Hierarchy	5
2.1.2 Key Performance Metrics	6
2.1.3 Data Characteristics and the Visualization Gap	7
2.2 Data Visualization	7
2.2.1 Visual Encoding and Graphical Perception	8
2.2.2 Chart Type Selection and Cognitive Constraints	9
2.3 Mobile Constraints and Design Patterns	10
2.3.1 Screen Real Estate and Information Density	10
2.3.2 Touch Interaction and Gesture Navigation	11
2.3.3 Progressive Disclosure and Layered Access	11
2.4 Usability and Information Clarity	12
2.4.1 Defining Usability in a Mobile Context	12
2.4.2 Information Clarity as Evaluation Dimension	13
2.4.3 Cognitive Load and Data Comprehension	14
3 Analysis and Design Space	15
3.1 Domain Requirements	15

3.2	Visualization Challenges on Mobile Devices	16
3.3	Derivation of Three Visualization Strategies	18
4	Visualization Strategies and Prototypes	20
4.1	Shared Data Basis and Metric Set	20
4.2	Strategy A: Card Based Progressive Disclosure	21
4.3	Strategy B: Narrative Contextualization with Category Browsing	24
4.4	Strategy C: Multi Chart Exploration with Carousel Navigation	26
4.5	Implementation and Prototype Fidelity	28
5	Evaluation Design	31
5.1	Study Framework	31
5.2	Participants	32
5.3	Procedure	33
5.4	Task Scenarios	33
5.5	Measurement Instruments	35
5.5.1	Task Performance and Experience Items	35
5.5.2	Modified System Usability Scale	36
5.5.3	Comparative Rating and Final Assessment	36
6	Results	37
6.1	Task Completion and Efficiency	37
6.2	Usability Scale Comparison	39
6.3	Information Clarity Across Strategies	41
6.4	Comparative Ranking and Preference	44
7	Discussion	47
7.1	Interpretation of Comparative Findings	47
7.2	Trade-offs Between the Strategies	49
7.3	Implications for Mobile Data Visualization	51
7.4	Limitations and Threats to Validity	53
8	Conclusion and Outlook	56
	Bibliography	XII

Appendix	XVI
A Task Scenario Descriptions	XVI
B Questionnaire	XVII
C Raw Evaluation Data	XXI
D Prototype Screenshots	XXVI
E Mock Data Generation Parameters	XXXI
F Contents of Digital Media	XXXII
Eidesstattliche Erklärung	XXXIII

List of Figures

2.1	Perceptual ranking of elementary graphical tasks based on judgment accuracy, adapted from Cleveland and McGill [1].	8
2.2	Relationship between evaluation constructs and their measurement instruments.	13
3.1	Three-level data hierarchy underlying all three visualization strategies.	17
4.1	Home screen comparison: Strategy A native iOS (left), Strategy B Figma (center), Strategy C Figma (right).	21
4.2	Strategy A: Home with stacked metric cards (left), Categories donut chart (center), bottom sheet filters (right).	22
4.3	Strategy A detail views: Transactions stacked area (left), Transaction Fees stacked bar (center), Gas Fees ranked bar (right).	23
4.4	Strategy B: dApp detail (left), Home with narrative text (center), Categories list (right).	24
4.5	Strategy B metrics: Transactions (left), Wallets/UAW (center), Inflow/Outflow (right).	25
4.6	Strategy C: Home with KPI grid (left), Categories icon grid (center), Categories allocation (right).	26
4.7	Strategy C chart carousel: stacked bar (left), line chart (center), radial bar (right).	27
4.8	MVVM architecture of the Strategy A native iOS application.	28
4.9	Tooltip fidelity: native drag gesture (Strategy A, left) vs. static Figma tooltips (Strategies B and C).	29
5.1	Distribution of self-reported blockchain and crypto knowledge across the ten participants.	32
5.2	Evaluation structure mapping task scenarios to prototype assignments, step sequences, and target metrics.	34
6.1	Distribution of overall prototype ratings per strategy.	38
6.2	Mean ratings per modified SUS item across the three strategies.	40

6.3	Comparative ratings for data clarity and ease of use per strategy. Strategy C shows the largest divergence between the two dimensions.	44
6.4	Overall ranking distribution across the three strategies (n = 10).	46
8.2	Strategy A detail views. Left: UAW detail sheet with tooltip and stacked bar chart. Center: Time scale filter with aggregation mode, time range presets, and date picker. Right: NFT category detail with ranked project table. . . .	XXVI
8.3	Strategy B overview. Top row: Home screen with price chart, KPI row, and narrative text; Categories list; Core Metrics overlay with tooltip. Bottom row: Outmine dApp detail; Core Metrics wallet breakdown; Outmine UAW line chart with tooltip.	XXVII
8.4	Strategy B detail views. Left: Gaming category transaction chart carousel. Center: Search overlay with recently searched items. Right: Search with keyboard input filtering results.	XXVIII
8.1	Strategy A overview. Top row: Home screen with stacked metric cards, Categories tab with donut chart, time scale filter sheet. Bottom row: Transactions detail (stacked area), Transaction Fees detail (stacked bar), Gas Fees detail (ranked horizontal bar).	XXIX
8.5	Strategy C overview. Top row: Home screen with KPI grid; Categories icon grid; Categories ranked list with donut chart. Bottom row: Gaming Transactions (stacked area); Gaming UAW (stacked bar); Transaction Fees (radial bar chart).	XXX
8.6	Strategy C detail views. Left: Best performance treemap for Gaming wallet data. Center: Total UAW overview for Gaming. Right: Outmine Inflow/Outflow bidirectional bar chart.	XXXI

List of Tables

3.1	Comparison of the three visualization strategies across key design dimensions.	19
4.1	Prototype fidelity across implementation dimensions.	30
5.1	Task scenario assignment across prototypes.	35
6.1	Task completion status per strategy.	37
6.2	Overall prototype ratings (1 to 10 scale).	38
6.3	Modified SUS composite scores per strategy.	39
6.4	Item-level means across the five modified SUS items. Higher values indicate stronger agreement; for negative items, lower scores are preferable.	39
6.5	Modified SUS composite scores by expertise subgroup.	41
6.6	Per-task experience item means (1 = Strongly Disagree, 5 = Strongly Agree).	42
6.7	Cross-strategy readability (E1) by participant. Cells marked with – indicate missing data.	43
6.8	Part D comparative ratings (1 = Poor, 5 = Excellent).	43
6.9	Context-specific preference votes (n = 10).	45
6.10	Overall ranking distribution (n = 10).	45
6.11	Recommendation score distribution (1 to 10 scale, n = 10).	45
8.1	Part A: Participant background items.	XVIII
8.2	Part B, Task 1 (Strategy A): Experience items for interactive chart exploration.	XIX
8.3	Part B, Task 2 (Strategy B): Experience items for search-based discovery. . .	XIX
8.4	Part B, Task 3 (Strategy C): Experience items for report-based scanning. . .	XX
8.5	Part C: Modified SUS items (5-item variant). Scoring: positive items contribute (response – 1), negative items contribute (5 – response). The sum is multiplied by 5 to yield a composite score between 0 and 100.	XX
8.6	Part D: Cross-prototype comparative rating dimensions.	XXI
8.7	Part D: Context-specific preference items.	XXI
8.8	Part E: Final assessment and ranking items.	XXI
8.9	Participant demographics.	XXII

8.10	Task completion status and overall prototype ratings (1–10) per participant.	XXII
8.11	Experience items (1–5 Likert) per participant across all strategies. Strategy A: E1 Readable, E2 Comprehension, E3 Findability, E4 Information amount, E5 Trend visibility. Strategy B: E1 Readable, E2 Comprehension, E3 Findability, E4 Search quality. Strategy C: E1 Readable, E2 Overview, E3 Scannability, E4 Findability.	XXIII
8.12	Modified SUS item responses (S1–S5) and composite scores (Σ) per participant. S1: Overview (+), S2: Complexity (–), S3: Organization (+), S4: Clutter (–), S5: Confidence (+). Scoring: positive items contribute (response – 1), negative items contribute (5 – response), sum $\times$ 5.	XXIV
8.13	Part D comparative ratings: data clarity and ease of use (1–5) per participant.	XXIV
8.14	Overall ranking and recommendation score (1–10) per participant.	XXV
8.15	Contents of the digital media.	XXXII

List of Abbreviations

API	Application Programming Interface
dApp	Decentralized Application
DeFi	Decentralized Finance
iOS	iPhone Operating System
KPI	Key Performance Indicator
MVVM	Model-View-ViewModel
NFT	Non-Fungible Token
NPS	Net Promoter Score
SUS	System Usability Scale
TVL	Total Value Locked
UAW	Unique Active Wallets
UI	User Interface
UX	User Experience
WCAG	Web Content Accessibility Guidelines

1 Introduction

1.1 Problem Context and Motivation

Blockchain ecosystems generate performance data at multiple levels of aggregation simultaneously. A single network, such as Moonbeam, one of dozens in the Polkadot ecosystem, hosts applications across categories like DeFi, Gaming, and Infrastructure, each tracked through indicators including transaction volume, unique active wallets, gas fees, and asset price [2]. Because the data is temporal, categorical, and hierarchical simultaneously, presenting it on a mobile screen is not simply a matter of selecting the right chart type. It requires a visualization strategy: a coherent set of decisions about layout, navigation structure, information density, and interaction patterns that together determine how a user encounters and moves through the data.

Desktop tools for this purpose exist, but they serve a different paradigm. Tovanich et al. [3] found that blockchain visualization overwhelmingly targets large screens with pointer-based interaction, focusing on transaction-level studies rather than ecosystem-wide overviews. A companion review of Bitcoin visualizations confirms this orientation [4]. Foundation teams managing Layer-1 and Layer-2 networks increasingly need to access the same data on mobile devices: during conferences, in stakeholder meetings, or while reviewing ecosystem health outside of office settings. The desktop paradigm does not transfer to a phone without a fundamental redesign.

Mobile devices constrain visualization in ways that go beyond reduced screen size. Touch input replaces the precision of a mouse cursor. Hover interactions have no native equivalent on a touchscreen. The physical display area of a smartphone measures roughly six inches diagonally [5], forcing designers to choose between showing more data at lower legibility or less data at higher clarity. These constraints create trade-offs that have no single correct resolution. Usability and information clarity can diverge: an interface may be navigable in the sense that users find what they are looking for, yet unclear in the sense that they misread what they find [6]. Recognizing this distinction motivates a comparative approach.

Despite the practical relevance of this intersection, the research landscape shows a gap on

both sides. Blumenstein et al. [5] note that empirical work on mobile data visualization remains scarce, mainly regarding end-user requirements and actual usage contexts. Blockchain visualization research concentrates on expert-oriented desktop environments and rarely includes formal evaluation of the user experience [7]. No existing study compares alternative visualization strategies for blockchain ecosystem data on mobile devices, and no empirical evidence exists on whether strategy-level design choices affect usability and information clarity differently.

FiDi Tech Inc., a San Francisco-based company specializing in blockchain analytics, provides the industry context for this thesis. The company’s existing desktop platform serves foundation teams at organizations such as Moonbeam, zkSync and Mantle. A preceding project phase produced three conceptually distinct visualization strategies for a mobile version of this platform [2]. One organizes the data behind nested cards that reveal detail on demand. Another pairs charts with narrative text that guides interpretation. A third distributes multiple chart types across swipeable panels for user-driven exploration. All three share the same underlying data but differ in how they balance density against clarity, and how much interpretive effort they leave to the user. This thesis investigates whether those differences produce measurable outcomes.

1.2 Research Question and Objectives

The primary research question is:

How do different visualization strategies for blockchain network data perform on mobile devices in terms of usability and information clarity?

The investigation proceeds from design through implementation to evaluation. Three visualization strategies address the same blockchain data domain but diverge in layout, interaction pattern, and information architecture. Each is realized as a testable prototype at sufficient fidelity for comparative assessment. A within-subjects evaluation with participants ranging from blockchain analytics professionals to users with no prior exposure to technical applications provides the empirical basis. From these comparisons, the thesis derives findings on which design patterns support usability, which support information clarity, and where the two dimensions converge or conflict.

Usability, following ISO 9241-11 [8], encompasses effectiveness, efficiency, and satisfaction. Information clarity, as operationalized in this thesis, refers to the degree to which a visu-

alization allows its viewer to correctly extract, interpret, and compare the encoded data. Section 2.4 defines both constructs in detail. Measuring them separately reveals whether the mobile trade-offs force compromises that benefit one dimension at the expense of the other.

1.3 Methodological Overview

The thesis follows a design-oriented research approach that moves through four phases. A literature review establishes the theoretical foundations: visual encoding principles, mobile visualization constraints, cognitive load theory, and usability measurement frameworks. These foundations inform the second phase, in which domain requirements from the preceding project phase are mapped against mobile design tensions to derive three visualization strategies. The third phase implements these strategies as testable prototypes: one as a native iOS application built with Swift and SwiftUI, and two as interactive Figma prototypes. The fourth phase conducts a comparative evaluation with ten participants using a within-subjects design, collecting task performance data, modified System Usability Scale scores, information clarity ratings, and comparative preference rankings.

The evaluation instruments were administered remotely through Google Forms. Participants interacted with each prototype on their personal mobile devices and rated it across multiple quality dimensions. The asymmetry between the native application and the Figma prototypes introduces a fidelity gap that the evaluation design addresses by separating visual and structural assessment from interaction quality assessment.

The fidelity asymmetry between the native application and the Figma prototypes is a deliberate trade-off: it introduces a confound that the evaluation cannot resolve, but it also enables the comparison of interaction qualities that a uniform Figma implementation would have suppressed entirely.

1.4 Thesis Structure

The theoretical groundwork occupies Chapter 2: blockchain network data, visual encoding, cognitive load, mobile constraints, and the two evaluation constructs. Chapter 3 applies these foundations to the design problem by documenting domain requirements, identifying three visualization challenges, and deriving three strategies that address them differently.

Concrete prototypes are presented in Chapter 4, covering screens, chart types, navigation structures, the native iOS architecture, and a fidelity comparison across implementations. The evaluation design in Chapter 5 specifies the within-subjects framework, participant sample, task scenarios, and a five-part questionnaire.

Results appear in Chapter 6 without interpretation. Chapter 7 provides that interpretation, examining trade-offs between strategies, implications beyond the blockchain domain, and threats to validity. A direct answer to the research question, a summary of contributions, and directions for future work conclude the thesis in Chapter 8.

2 Foundations

2.1 Blockchain Network Data

A blockchain operates as a shared transaction record maintained collectively by a network of nodes, none of which holds singular control over the data [9]. Early systems like Bitcoin functioned as single, self-contained networks. Current architectures have moved beyond this model. Ecosystems such as Polkadot connect dozens of specialized blockchains into a shared infrastructure, each hosting its own set of decentralized applications, commonly referred to as dApps. The result is a layered data landscape: activity occurs at the level of individual applications, within specific networks, and across the ecosystem as a whole.

2.1.1 Network Architecture and Data Hierarchy

Wood [10] designed Polkadot as a heterogeneous multi-chain framework. A central Relay Chain provides shared security, while connected parachains operate as independent blockchains with their own state transition logic [11]. The Web3 Foundation describes interoperability between these chains as a core design goal [12]. This thesis uses the Polkadot ecosystem as its reference case because its multi-chain structure produces data at several distinct levels of aggregation, which makes the visualization problem particularly tangible.

Each network hosts dApps whose purpose and behavior vary widely. A DeFi lending protocol generates different activity patterns than a gaming application or an infrastructure service [11]. Labels such as DeFi, Gaming, or Social are not technical designations but analytical groupings, formed after the fact to make large application landscapes navigable. Tovanich et al. [3] reflect this in their systematic review by organizing blockchain visualization tools along data types and application domains rather than along protocol-defined boundaries.

The data hierarchy that results from this structure, ecosystem to network to category to dApp, is therefore not built into the blockchain architecture. It is constructed as an organizational strategy for visualization and monitoring. Sarikaya et al. [13] observe that this is not unusual: in dashboard design more broadly, the structure of a display tends to

reflect analytical intent rather than the raw topology of the underlying data. For blockchain ecosystems, any dashboard must first impose this constructed hierarchy before it can present activity in a readable way.

2.1.2 Key Performance Metrics

Five quantitative indicators form the core of ecosystem performance monitoring: Unique Active Wallets (UAW), transaction count, transaction fees, gas fees, and asset price.

UAW measures how many distinct wallet addresses interact with a network or dApp within a given time window. The concept parallels established Web2 engagement metrics such as Daily Active Users [14]. A wallet counts as active if it has executed at least one meaningful on-chain action in the specified period. The metric has a well-known limitation: wallet addresses are pseudonymous, and a single person can control multiple wallets [9]. UAW is therefore most accurately described as a relative signal for comparing adoption throughout dApps or networks, not as an exact count of real users.

Transaction count and fee metrics cover the operational side. Transaction volume shows how much activity a network processes. Gas fees describe the computational cost of a single operation, while transaction fees represent the total cost borne by the user. Ferretti and D'Angelo [9] show that periods of high demand manifest as spikes in transaction counts. High transaction numbers combined with low fees suggest an efficient, actively used chain. The reverse may indicate congestion or limited capacity. Asset price tracks the market valuation of a network's native token, introducing a volatile, externally driven variable whose interpretive function is contextual: wallet and transaction metrics gain meaning when considered alongside price dynamics.

All five metrics must be tracked for every dApp within every network. A single network hosting twenty dApps already produces one hundred data points per time interval, before any temporal dimension is considered. The visualization challenge is not to display one metric for one entity, but to present five metrics across dozens of entities in a way that supports comparison, trend recognition, and categorical grouping on a mobile phone screen.

2.1.3 Data Characteristics and the Visualization Gap

Shneiderman [15] distinguished several data types that each require their own visual strategies: temporal, multidimensional, hierarchical, and network data. Blockchain performance metrics do not fall neatly into one of these categories. UAW and transaction counts tracked over weeks are temporal. Comparing gas fees across dApps within a category is categorical and multidimensional. The layered organization from ecosystem down to dApp is hierarchical. The practical consequence is that no single chart type can represent the full picture.

Tovanich et al. [3] surveyed 76 blockchain visualization tools and found that time series plots are the most common representation, followed by node-link diagrams for transaction networks. The tools they reviewed are predominantly built for desktop use and focus on individual addresses or transactions rather than aggregated ecosystem views. A companion review of Bitcoin visualizations confirms this orientation [4]. Representative systems such as BitExTract [16], HyperSec [17], and similar tools [18] target desktop environments with pointer-based interaction. Napoli et al. [7] observe that the field concentrates on expert-oriented environments and rarely includes formal evaluation of the user experience. Dashboards at the ecosystem level that aggregate activity across a multi-chain hierarchy are absent from the surveyed literature.

Blumenstein et al. [5] note a parallel gap in mobile visualization research: empirical work on end-user requirements and real-world usage contexts remains rare. Burigat and Chittaro [19] demonstrate that even well-established techniques like Overview+Detail need fundamental rethinking on mobile devices. The intersection of these two gaps, blockchain ecosystem data and mobile presentation, defines the problem space this thesis addresses.

2.2 Data Visualization

Every chart type, color mapping, and axis label used to present blockchain data encodes assumptions about what the viewer should notice first and what can be deferred. Translating the dataset into something a person can read on a screen requires encoding decisions with real consequences for accuracy and cognitive effort, consequences that intensify when multiple chart types compete for limited screen space.

2.2.1 Visual Encoding and Graphical Perception

Cleveland and McGill [1] tested how accurately subjects could judge proportional differences when values were embedded in different chart formats. Their experiments produced a ranked list of elementary perceptual tasks ordered by judgment precision. Position along a common scale yields the most accurate readings. Length without a shared baseline is less precise. Angle and area fall further behind, and color saturation ranks near the bottom for quantitative comparisons.




Figure 2.1: Perceptual ranking of elementary graphical tasks based on judgment accuracy, adapted from Cleveland and McGill [1].

The ranking has direct consequences for chart selection. A bar chart encodes values as positions on a shared axis, placing it at the top of the accuracy hierarchy. A pie chart relies on angle and area, both judged with considerably more error. Heer and Bostock [20] replicated these experiments using Amazon Mechanical Turk with a larger and more diverse subject pool; the original ranking held. Their crowdsourced setting also surfaced practical thresholds: chart height showed diminishing returns for accuracy beyond approximately 80 pixels, and gridlines lost effectiveness below roughly 8 pixels of spacing. On a mobile dashboard, where each chart competes for vertical space, these thresholds matter immediately.

Munzner [21] systematized these findings through a marks and channels framework. Two principles govern assignment: expressiveness requires that a channel's perceptual properties match the data type, and effectiveness requires that the most important attributes receive the most accurate channels. For a blockchain dashboard, this means encoding high-priority quantitative attributes like transaction volume through position on a common scale rather than through area or color intensity.

2.2.2 Chart Type Selection and Cognitive Constraints

The encoding hierarchy narrows the design space but does not resolve the selection problem entirely. Munzner [21] offers a guideline: if the key attribute is temporal, use a line chart; if categorical, use a bar chart. For part-to-whole relationships, stacked bar charts and stacked area charts are standard options. Donut and pie charts serve a similar function but rely on angle and area encoding, placing them lower in the perceptual hierarchy. They remain common in dashboard design because they communicate proportional structure at a glance, even if precise value comparison suffers.

The blockchain data described in Section 2.1 requires several chart types at once. Temporal trends suit line charts. Comparing gas fees across dApps calls for bar charts. Showing category contributions to total activity requires stacked or proportional representations. A single dashboard screen must contain these different encodings side by side, each calling for its own axis, legend, and spatial allocation.

At this density, the limits of human information processing become a design constraint. Sweller [6] introduced Cognitive Load Theory to describe how working memory capacity constrains learning and problem solving. Working memory can hold and process only a limited number of information elements at any given time. Sweller distinguished between intrinsic load, determined by the complexity of the material itself, and extraneous load, caused by the way the material is presented. A confusing layout or an unintuitive color scheme generates extraneous load without contributing informational value. Section 7.2 revisits this distinction when examining how progressive disclosure distributes interpretive effort across navigation layers.

Applied to visualization, position encoding along a clear axis imposes relatively low extraneous load. Area encoding with ambiguous labels forces the viewer to work harder for the same information. When multiple charts appear on one screen, this effect accumulates: each additional legend, axis, and color mapping requires more contextual information in work-

ing memory. Sarikaya et al. [13] found that dashboard designers frequently struggle with this accumulation, adding visual elements that increase informational coverage while risking overwhelming the viewer. Beyond a certain threshold, simultaneously visible encodings exceed what a user can process, and additional information becomes noise rather than signal. The visualization strategies compared in this thesis differ in how they manage this tension.

2.3 Mobile Constraints and Design Patterns

Visualization research has historically assumed a desktop environment: a large monitor, a mouse for exact pointing, and enough screen space to show overview and detail side by side. Mobile devices break each of these assumptions. The screen is smaller by an order of magnitude, the input mechanism is a fingertip rather than a cursor, and the viewing context shifts from a stationary desk to a moving hand.

2.3.1 Screen Real Estate and Information Density

Modern smartphones occupy a diagonal range of roughly 4.7 to 6.7 inches [5]. A chart that fills half a desktop monitor might be allotted two inches of vertical space on a phone. Pixel density compensates partially: text and thin lines can remain sharp on displays exceeding 400 pixels per inch. Physical area, however, remains fixed, and the user's finger still covers roughly the same fraction of the screen regardless of resolution.

The central tension is between information density and readability. Fitting a blockchain dashboard onto a phone without interaction means compressing charts, shrinking labels, and eliminating whitespace. At some point, the compression undermines the encoding accuracy established by the perceptual thresholds in Section 2.2.1. Hoffswell et al. [22] studied how news organizations adapt visualizations to smaller screens and identified four responsive strategies: resizing, repositioning, transposing, and removing content. Removal was the most frequent. The strategy is practical but carries a cost: every removed element is information the viewer no longer has access to. Horak et al. [23] propose a more gradual approach with responsive levels of detail, defining four rendering stages tied to available cell size, from pixel-level abstraction to a medium view with axes and annotations. The visualization adapts its complexity to the space it receives, preserving at least a coarse signal where a binary removal strategy would show nothing.

Brehmer et al. [24] compared animation and small multiples for trend visualization on mobile phones. Small multiples led to faster responses on most tasks despite giving each individual frame less room, though accuracy varied by task type. Animation produced higher subjective confidence but did not improve task performance. The result suggests that on small screens, spatial juxtaposition outperforms sequential presentation, even when each juxtaposed view is itself quite small.

2.3.2 Touch Interaction and Gesture Navigation

Brehmer et al.¹ [25] define the fat-finger problem as the mismatch between the size of a human fingertip and the size of graphical marks in a visualization. Selecting a data point precisely becomes difficult, and the likelihood of accidentally activating a neighboring element rises. The problem is compounded by occlusion: a finger physically covers the area it is trying to interact with.

A more fundamental gap is the absence of hover. On a desktop, moving the mouse over a chart element triggers a tooltip at no interaction cost. Touch screens have no equivalent passive state. Snyder et al. [26] propose a drag-based inspect gesture as a replacement: the user places a finger near the area of interest and drags across the visualization, with a tooltip following the finger position. The gesture separates inspection from selection and must account for the user's thumb moving upward, as this may obscure labels at the top of the view. For a mobile blockchain dashboard, these constraints shape every interactive element: tooltips must respond to drag, time-range selectors must accommodate thumb operation, and metric cards need touch targets large enough to avoid misfires.

2.3.3 Progressive Disclosure and Layered Access

Shneiderman [15] formulated the Visual Information Seeking Mantra: overview first, zoom and filter, then details on demand. On a mobile screen, this assumption is strained. Burigat and Chittaro [19] investigated Overview+Detail interfaces on handheld devices and found that the two views compete for space the device does not have: making the overview readable steals space from the detail view, and vice versa.

Progressive disclosure offers a way out. Rather than showing overview and detail simulta-

¹Matthew Brehmer, Bongshin Lee, John Stasko, and Christian Tominski, "Interacting with Visualization on Mobile Devices," in *Mobile Data Visualization*, Chapman and Hall/CRC, 2021, pp. 67–110.

neously, the interface reveals information in layers. The first screen presents a reduced set of high-level indicators; interaction exposes the next level. Cognitively, each layer keeps the number of simultaneously visible elements within working memory limits [6], and transitions are driven by user intent rather than a static layout demanding parallel processing.

In mobile interface design, this principle materializes through card-based layouts, bottom sheets, and expandable sections, all of which defer detail until the user requests it. The visualization strategies compared in this thesis implement progressive disclosure at different granularities: a nested card layout stacks layers explicitly, while a report-style layout places more information on the first screen and relies on scrolling. Both manage the tension between overview and detail, but they distribute cognitive load differently, a distinction the empirical evaluation in Chapter 5 is designed to measure.

2.4 Usability and Information Clarity

Evaluating the result of applying visualization to mobile constraints requires two complementary perspectives: usability, concerned with whether users can operate the interface to accomplish their tasks, and information clarity, concerned with whether they actually understand the data it presents.

2.4.1 Defining Usability in a Mobile Context

ISO 9241-11 grounds usability in three qualities: effectiveness (whether users accomplish goals correctly and completely), efficiency (the resources spent), and satisfaction (perceived ease and comfort) [8]. None can be assessed without specifying who the users are, what they are trying to do, and under what conditions. Kurosu and Hashizume [27] sharpen the boundary by separating usability from user experience: usability is task-oriented, while user experience encompasses emotional and esthetic responses. A comparative study benefits from the task-oriented lens because it grounds measurement in observable outcomes.

Mobile use adds contextual factors: fragmented attention, limited screen space, and imprecise input. Blumenstein et al. [5] observe that studies rarely report the physical usage environment, even though distractions and posture affect performance. Usability scores obtained in a controlled setting therefore reflect interface quality under favorable conditions.

2.4.2 Information Clarity as Evaluation Dimension

Information clarity, as used in this thesis, refers to the degree to which a visualization allows its viewer to correctly extract, interpret, and compare the encoded data. Borgo et al. [28] summarize three levels of visualization understanding originally proposed by Bertin. Elementary reading extracts a single data point. Intermediate reading recognizes trends across multiple data points. Comprehensive reading compares structures across views and requires domain knowledge. A blockchain dashboard that performs well in terms of information clarity supports all three levels.




Figure 2.2: Relationship between evaluation constructs and their measurement instruments.

Clarity can be measured through two paths. Task-based accuracy tests the elementary and intermediate levels through verifiable questions about extracted data values or recognized trends. Subjective instruments capture perceived clarity through self-report items without requiring ground-truth verification. The evaluation in this thesis follows the subjective path, using per-task experience ratings and a cross-prototype clarity comparison rather than value-extraction tasks with verifiable answers. This choice limits the findings to perceived rather than verified clarity, a constraint discussed in Section 7.4. Sarikaya et al. [13] note that perceived clarity tends to degrade as chart elements accumulate on a single screen, a prediction the three visualization strategies allow to test even through subjective measurement.

2.4.3 Cognitive Load and Data Comprehension

The cognitive load constraints introduced in Section 2.2.2 apply with particular force on mobile devices. The data complexity of a blockchain dashboard produces substantial intrinsic load, while mobile constraints amplify the extraneous component: compressed charts make labels harder to read, small touch targets increase interaction effort, and the absence of hover forces additional gestures. Munzner [21] frames this as a design obligation: effective visualization should amplify cognition, not merely display data. The empirical evaluation in this thesis examines whether the three compared strategies differ in the cognitive effort they impose.

3 Analysis and Design Space

Translating theoretical foundations into design decisions requires a concrete problem. The domain context provides it: blockchain ecosystem data with specific metrics, user groups, and monitoring tasks. From this context, three visualization tensions emerge where data complexity meets mobile constraints. The final section derives three visualization strategies whose prototypes Chapter 4 describes in detail.

3.1 Domain Requirements

The industry context introduced in Section 1.1 shapes the design requirements directly. FiDi Tech’s analytics platform currently serves foundation teams through desktop interfaces [2]. Foundation staff at organizations such as Moonbeam, zkSync, Mantle, Eigenlayer, and Optimism increasingly need mobile access to review ecosystem health outside of office settings, during conferences, stakeholder meetings, or between scheduled reviews. The preceding project phase translated this demand into concrete requirements.

Target Users and Core Tasks

The preceding project phase identified two user groups and five core use cases [2]. Analytics teams within blockchain foundations form the primary group. These users work regularly with metrics such as transaction volume, unique active wallets (UAW), gas fees, and total value locked (TVL). Their tasks range from ecosystem health monitoring to grant performance evaluation and cross-category comparisons for resource allocation. A secondary group consists of marketing and business development staff, who need quick performance overviews for investor updates, exportable visualizations for board presentations, and data summaries suited to community communications.

From these profiles, five core use cases emerged: retrieving an ecosystem performance overview that summarizes key metrics across the entire network, comparing grant performance of funded projects against one another, analyzing category-specific metrics where

categories such as DeFi, Gaming, or Infrastructure group related dApps, visualizing time series development at multiple granularity levels for governance decisions, and identifying top performing dApps for marketing materials.

Functional and Non-Functional Requirements

The prototypes must support metrics with distinct data characteristics: time-based metrics, particularly transaction volume and fee development; categorical metrics, such as dApp category distributions; and hierarchical metrics structured along a three-level hierarchy of ecosystem, category, and individual dApp. Navigation throughout all three levels must enable drill-down from the ecosystem overview through category aggregations to individual dApp detail views. Visualization forms must be selected in a context-aware manner, following assignment rules derived from the perceptual principles discussed in Section 2.2.1. To enable comparative evaluation, alternative interaction patterns for the same functionality must coexist within the prototypes.

Non-functional requirements complement these specifications. Load times must remain below three seconds and the frame rate must hold at 60 frames per second, preventing technical artifacts from distorting usability measurements. Accessibility requirements following WCAG 2.1 Level AA apply, including touch targets of at least 44 by 44 points, so that evaluation results reflect design quality rather than accessibility barriers [2].

3.2 Visualization Challenges on Mobile Devices

Where the domain requirements demand richness, the mobile constraints analyzed in Section 2.3 impose reduction. The intersection produces three design tensions that define the problem space for all three visualization strategies.

Hierarchical Depth versus Screen Constraints




Figure 3.1: Three-level data hierarchy underlying all three visualization strategies.

Blockchain ecosystem data is inherently hierarchical. A single network such as Moonbeam contains multiple dApp categories, each grouping dozens of individual applications, all tracked across five or more metrics over time. Displaying this structure on a screen measuring roughly 390 by 844 logical points forces decisions about what to show and what to defer. Hoffswell et al. [22] catalog the available responsive techniques: elements can be removed, resized, repositioned, or transposed. Removal dominates on mobile, meaning that every hierarchy level beyond the first must be accessed through interaction rather than simultaneous display. How many levels can coexist on a single screen before information density exceeds the readability thresholds discussed in Section 2.3.1 is a question each strategy answers differently.

Simultaneous Metrics versus Sequential Access

Foundation analysts compare multiple metrics to form a judgment. A grant evaluation might require simultaneous awareness of transaction count, wallet activity, and fee structure for a single dApp. On desktop, these metrics occupy adjacent panels. A mobile screen accommodating more than two or three chart elements produces the clutter that Sarikaya et al. [13] identify as a recurring dashboard problem. Sequential access, where the user switches between metric views through taps or swipes, preserves readability for each individual metric but shifts the burden to working memory. The user must retain the value from the previous

view while interpreting the current one. Sweller’s cognitive load framework [6] predicts that this retention task generates extraneous load competing with the analytical task itself.

Navigation Depth and the Overview Problem

The number of interactions required to reach a specific data point serves as a proxy for navigation complexity. In a three-level hierarchy with four metrics and multiple time granularities, the path from the home screen to a specific chart for a particular dApp can involve three to five taps. Each tap is a context switch that risks disorientation if the interface does not preserve sufficient context. Shneiderman’s mantra of overview first, zoom and filter, then details on demand [15] provides a guiding principle, but Burigat and Chittaro [19] found that overview plus detail layouts tend to fail on small screens. A home screen showing all core metrics provides overview but limits trend visibility. Adding sparklines improves trend awareness but may exceed the capacity thresholds described by Horak et al. [23]. The three strategies resolve this tension differently.

3.3 Derivation of Three Visualization Strategies

During the preceding project phase, three conceptually distinct UI/UX prototypes were designed [2]. All three operate on the same data, the Moonbeam ecosystem, using identical metrics. They diverge along three axes that emerged from the challenges above: information density, degree of user guidance, and chart variety. Table 3.1 summarizes these divergences; Chapter 4 presents the prototypes themselves.

Strategy A: Card Based Progressive Disclosure. This strategy addresses the hierarchy and screen tension through strict separation of overview and detail. The home screen limits itself to an asset price line chart and numeric KPI cards. Metric specific visualizations appear only on demand in dedicated detail views. Chart types remain within the position-based encodings that Cleveland and McGill [1] and Heer and Bostock [20] rank highest for perceptual accuracy. The approach trades trend visibility at the overview level for low visual density and high per-chart readability, consistent with the benefits Brehmer et al. [24] associate with reduced element counts on small screens.

Strategy B: Narrative Contextualization with Category Browsing. Textual interpretations embedded alongside visualizations distinguish this strategy. Where Strategy A

defers interpretation to the user, Strategy B provides guided scaffolding through narrative descriptions that accompany each chart. This offloads part of the interpretive effort from working memory, following the cognitive load reduction principles discussed in Section 2.4.3. The trade-off is that narrative text occupies screen space, increases density per viewport, and may constrain analytical autonomy by foregrounding specific readings of the data.

Strategy C: Multi Chart Exploration with Carousel Navigation. Rather than removing visual complexity, this strategy distributes it across swipeable panels, following the Levels of Detail framework by Horak et al. [23]. Multiple chart types for the same dataset let the user select the encoding that matches their analytical need. The risk is that swiping between representations fragments working memory: the previous chart disappears, forcing comparison to rely on recall rather than simultaneous perception.

Dimension	Strategy A	Strategy B	Strategy C
Information density	Low (price chart + KPI values on home)	Medium (price chart + KPIs + narrative text)	High (price chart + KPI grid + wallet visuals)
Navigation tabs	3 (Home, Categories, Search)	2 (Home, Categories)	5 (Home, Categories, Projects, Favorites, variable)
Detail access	Bottom sheet on tap	Inline overlay with filters	Carousel swipe + timeline slider
User guidance	Minimal (numeric cues)	High (narrative text)	Low (free exploration)
Chart variety	4 types	5 types	6 types
Time granularity	Full (1H to All + date picker)	Limited (month tabs + 30D)	Full (1H to All + timeline slider)
Implementation	Native iOS (Swift/SwiftUI)	Figma prototype	Figma prototype

Table 3.1: Comparison of the three visualization strategies across key design dimensions.

Strategy A’s native implementation and the Figma-based realization of Strategies B and C create a fidelity asymmetry whose implications Section 4.5 documents and Section 7.4 evaluates.

4 Visualization Strategies and Prototypes

Chapter 3 derived three visualization strategies from the domain requirements and the design tensions between blockchain data complexity and mobile constraints. This chapter describes how those strategies materialize as prototypes. A shared data basis underlies all three, ensuring that differences in evaluation outcomes reflect design choices rather than content variation. The subsequent sections walk through each prototype’s screens, chart types, and navigation structure. The chapter closes with a discussion of the fidelity gap between the native application and the two Figma prototypes.

4.1 Shared Data Basis and Metric Set

All three prototypes draw on the Moonbeam network as their ecosystem context. Moonbeam is a Layer-1 blockchain in the Polkadot ecosystem, selected because FiDi Tech’s existing analytics platform already tracks it across multiple metric types [2]. The shared metric set comprises the five core indicators introduced in Section 2.1.2: transaction volume, unique active wallets (UAW), gas fees, transaction fees, and asset price. Each metric can be disaggregated along the three-level hierarchy described in Section 3.1: ecosystem totals, category aggregations across twelve dApp categories, and individual dApp values for 17 tracked applications.

Strategy A generates its data through a mock data module with deterministic generator functions seeded with fixed values, producing structurally realistic series without requiring network connectivity. The Figma prototypes use static screen data derived from the same Moonbeam ecosystem structure, presenting equivalent category names, metric types, and value ranges. Section 4.5 describes the technical architecture in detail.




Figure 4.1: Home screen comparison: Strategy A native iOS (left), Strategy B Figma (center), Strategy C Figma (right).

4.2 Strategy A: Card Based Progressive Disclosure

The home screen is organized around two functional zones. The upper section displays an interactive asset price line chart with the current value and percentage change. Below it, four metric cards are stacked using a Z-axis offset pattern, each showing only the metric title, an icon, and a single aggregated value.

Tapping a card opens a full-screen bottom sheet containing the corresponding detail visualization. Each sheet follows a consistent template: a header row, the chart itself, a color-coded legend, and a time filter control. The chart type varies by metric. Transactions use a stacked area chart, UAW uses a stacked bar chart, Gas Fees use a horizontal ranked bar chart showing the top ten contributors, and Transaction Fees use a stacked bar chart. Touch interaction on charts follows a drag gesture pattern: moving a finger horizontally selects the nearest data point and displays a tooltip listing all visible series with their values for the selected date.

The Categories tab presents a donut chart where each segment represents a dApp category




Figure 4.2: Strategy A: Home with stacked metric cards (left), Categories donut chart (center), bottom sheet filters (right).

weighted by UAW. Tapping a segment updates the center overlay to show the selected category's totals. Below the donut chart, stacked category cards display category level aggregations, and tapping one opens a detail sheet with a ranked project table. The Search tab provides a text field with live filtering across all tracked tokens and dApps, linking to detail views that reuse the same dashboard layout scoped to the selected entity.

Navigation across all three tabs is managed through a shared state object that synchronizes the selected network, aggregation mode, and date range across views.




Figure 4.3: Strategy A detail views: Transactions stacked area (left), Transaction Fees stacked bar (center), Gas Fees ranked bar (right).

4.3 Strategy B: Narrative Contextualization with Category Browsing



Figure 4.4: Strategy B: dApp detail (left), Home with narrative text (center), Categories list (right).

Strategy B reduces the tab count to two: Home and Categories. The home screen opens with the asset price line chart, followed by a summary row showing current metric values alongside percentage changes. Monthly navigation tabs and a 30-day toggle control the time window. A narrative text block accompanies the primary visualization, stating in plain language which categories dominate and what directional trend the data exhibits.

The combination of chart and explanatory text is the distinguishing feature of this prototype. Rather than leaving the user to extract patterns independently, the narrative provides an explicit interpretation that the user can confirm or challenge against the visual evidence. Chip based filters above the chart area allow switching between metric types and selecting specific categories, and filters are combinable.

Category detail views present dApp lists with inline metrics that update when the active metric type changes. Chart types specific to this prototype include a treemap for wallet




Figure 4.5: Strategy B metrics: Transactions (left), Wallets/UAW (center), Inflow/Outflow (right).

distribution across categories and inflow/outflow bar charts showing bidirectional wallet movement. Time navigation is the most restricted of the three prototypes: monthly tabs and the 30 day toggle fit the report-oriented character, assuming the user desires to review a recent period rather than explore arbitrary date ranges.

4.4 Strategy C: Multi Chart Exploration with Carousel Navigation




Figure 4.6: Strategy C: Home with KPI grid (left), Categories icon grid (center), Categories allocation (right).

Strategy C uses a five-tab navigation: Home, Categories, Projects, Favorites, and a search function at the top. The home screen packs the highest information density of the three prototypes. An asset price line chart occupies the top area, and below it a KPI grid displays four metrics with current values and percentage changes alongside small circular visualizations for wallet metrics.

The Categories tab presents a large donut chart segmented by dApp category, with aggregate values displayed in the center. Two view modes are available for the category list: an icon grid and a ranked list showing transaction and UAW columns side by side. This dual view is unique to Strategy C and lets the user switch between a visual overview and a tabular comparison without navigating to a new screen.

Within a category detail view, the chart carousel is the central interaction element. A horizontal swipe moves between three representations of the same dataset: a stacked area



Figure 4.7: Strategy C chart carousel: stacked bar (left), line chart (center), radial bar (right).

chart foregrounding temporal trends, a stacked bar chart enabling part-to-whole comparisons at discrete time points, and a radial bar chart compressing categorical proportions into a compact circular layout. The carousel makes chart type selection an explicit user action rather than a design decision fixed by the prototype.

Time navigation is the most elaborate of the three designs. Inline buttons offer predefined scales from one hour to all time, and a draggable timeline slider allows custom date range selection. Together, these controls cover both quick preset access and fine grained temporal exploration.

4.5 Implementation and Prototype Fidelity

Technical Architecture of Strategy A

Strategy A was implemented as a native iOS application targeting iOS 17 and above, built with Swift and SwiftUI. The architecture follows the MVVM (Model-View-ViewModel) pattern. Three primary ViewModels manage the application state: **AnalyticsViewModel** for the dashboard, **CategoriesViewModel** for the categories tab, and **SearchViewModel** for the search functionality. All three use the `@Observable` macro introduced in iOS 17 for reactive state propagation.




Figure 4.8: MVVM architecture of the Strategy A native iOS application.

Data access is abstracted through a repository protocol whose mock implementation generates all data locally. Asset prices follow geometric Brownian motion with annual seasonality and mean reversion; metric series combine sinusoidal base patterns with randomized noise. The generation parameters are documented in Appendix 8. Charts are rendered using Apple’s Swift Charts framework with drag-based data point selection and categorical foreground styling. The application was distributed to evaluation participants via Apple TestFlight.

Figma Prototypes for Strategies B and C

Strategies B and C were designed as interactive Figma prototypes. Both define screen sequences with tap and swipe transitions between views. Chart tooltips, filter state changes,

and metric value updates are represented through linked screen variants rather than dynamic computation. The prototypes reproduce layout, typography, spacing, and color schemes at full visual fidelity. What they cannot replicate is the behavioral layer: drag gestures that continuously update a tooltip position, pull to refresh mechanics, haptic feedback, or real time data recomputation when filters change.

Fidelity Comparison




Figure 4.9: Tooltip fidelity: native drag gesture (Strategy A, left) vs. static Figma tooltips (Strategies B and C).

The fidelity gap means that Strategy A is evaluated under conditions closer to a production application, while Strategies B and C are evaluated as design prototypes. Research on prototype fidelity suggests that low-fidelity prototypes identify similar usability problems as high-fidelity versions, though task completion times measured on lower-fidelity artifacts tend to be unreliable predictors of final product performance [29]. This asymmetry is a deliberate methodological trade-off documented in the project report [2]: implementing all three strate-

gies as native applications was infeasible within the project schedule, but limiting all three to Figma would have prevented evaluation of the interaction qualities that differentiate mobile experiences from static mockups. The evaluation design described in Chapter 5 accounts for this asymmetry by separating assessment of visual and structural qualities, which all three prototypes support, from assessment of interaction qualities, where only Strategy A provides valid data.

Table 4.1: Prototype fidelity across implementation dimensions.

Dimension	Strategy A	Strategy B	Strategy C
Platform	Native (Swift/SwiftUI)	iOS Figma interactive prototype	Figma interactive prototype
Data	Dynamic (mock generation)	Static (screen variants)	Static (screen variants)
Chart interaction	Continuous drag + tooltip	Tap to show preset tooltip	Tap to show preset tooltip
Gestures	Native (swipe, drag, pull)	Simulated (tap, screen link)	Simulated (tap, screen link)
Haptic feedback	Yes (tab switch, refresh)	No	No
Loading states	Skeleton animation	Not represented	Not represented
Distribution	Apple TestFlight	Figma app / mobile browser	Figma app / mobile browser

5 Evaluation Design

The three visualization strategies described in Chapter 4 differ in information density, navigation structure, and chart variety, but they operate on the same data and serve the same user goals. Whether these differences translate into measurable outcomes for usability and information clarity is an empirical question.

5.1 Study Framework

A within-subjects comparative design structures the evaluation [30]. Every participant interacts with all three prototypes and rates each one individually, allowing for direct comparison across strategies without the confounds that arise when different groups evaluate different systems. The visualization strategy (A, B, or C) serves as the independent variable. Dependent variables fall into four categories: task performance measured through completion status, perceived usability captured by a modified System Usability Scale, information clarity assessed through per-task Likert items, and direct comparative preference ratings including data clarity and ease of use.

The study is exploratory rather than confirmatory. Ten participants drawn from a specialized domain do not support inferential statistics or hypothesis testing. The goal is to identify patterns, tendencies, and qualitative contrasts. Blumenstein et al. [5] report a median sample size of twelve across 21 mobile visualization studies, with lower-quartile values around six. Blattgerste et al. [31] note that SUS scores reach approximately 80 percent conclusiveness at ten participants, suggesting that samples of this magnitude yield interpretable usability data.

The fidelity asymmetry between prototypes (Section 4.5) constrains the task design. Tasks were assigned to the prototype whose interaction capabilities they require, rather than distributing all tasks across all prototypes. Crossing tasks with strategies would have required either restricting all three to Figma fidelity, eliminating the interaction dimension from comparison entirely, or assigning interaction-dependent tasks to prototypes incapable of supporting them. The modified SUS and Part D instruments compensate at the strategy level by

collecting ratings from every participant on every prototype, independent of task assignment.

5.2 Participants

Ten participants completed the study. Recruitment followed a purposive sampling strategy targeting a mix of domain experts and non-specialist users [30]. The sample is not limited to a single organization; participants span the broader Web3 ecosystem and adjacent technology sectors, including founders, developers, data analysts, and students at different organizations.




Figure 5.1: Distribution of self-reported blockchain and crypto knowledge across the ten participants.

Five participants work professionally with Web3 data and blockchain analytics, using tools such as DappRadar, DefiLlama, Dune, or custom dashboards on a regular basis. Three bring adjacent technical backgrounds without specific Web3 experience. Two reported no prior exposure to technical applications. Including participants across this range is deliberate: visualization strategies that communicate data structure and meaning only to domain experts would score well on information clarity in an expert sample yet fail the broader audience that mobile applications typically serve.

Three experience dimensions were collected through the background questionnaire: frequency of using visualization apps or dashboards, self-rated blockchain and crypto knowledge, and frequency of using mobile apps for financial or analytics data. An additional binary item captured whether participants had prior experience with dedicated blockchain analytics tools.

5.3 Procedure

The study was conducted remotely. Each participant received access to three separate Google Forms, one per prototype, alongside installation instructions. Strategy A was distributed via Apple TestFlight. Strategies B and C were accessed as interactive Figma prototypes through the Figma mobile app or a mobile browser.

Participants completed the evaluation at their own pace without a facilitator present. The Google Forms guided them through task steps, provided rating scales immediately after each task, and collected the SUS items and comparative ratings in a fixed sequence. The three forms could be completed in any order. Conducting the study remotely reflects the geographic distribution of the participant pool but introduces limitations: no think-aloud data or behavioral observations were captured, and participants who encountered difficulties had no immediate recourse.

5.4 Task Scenarios

Three task scenarios structure the evaluation. Each targets a specific interaction pattern and is assigned to the prototype that supports it at the required fidelity level.

Task 1: Interactive Chart Exploration (Strategy A). The participant uses the search function to find “Beamswap,” opens the detail page, locates the UAW chart, changes the time range to one year, and explores the chart interactively. This is the only scenario requiring continuous gesture input and dynamic system feedback.

Task 2: Search and Discover dApps (Strategy B). The participant navigates to the Categories tab, uses the search function to find “Outmine,” and locates the Inflow/Outflow visualization on the detail page. The task tests search-based discovery and the readability of a bidirectional metric visualization.

Task 3: Report-Based Overview (Strategy C). The participant navigates to the Categories tab, selects Gaming, scrolls through the report-style layout, and locates the Transaction Fees section. The task evaluates whether a scroll-based layout supports efficient scanning.




Figure 5.2: Evaluation structure mapping task scenarios to prototype assignments, step sequences, and target metrics.

Each prototype receives one task targeting its defining interaction paradigm. Assigning the interactive exploration task to a Figma prototype would test the limitations of the prototyping tool rather than the visualization strategy. Per-task comparisons across strategies are not possible; the cross-prototype instruments described below compensate for this. Table 5.1 summarizes the assignment structure. Full task descriptions appear in Appendix 8.

Task	Scenario	Prototype	Target Metric(s)	Met-	Interaction Focus
1	Interactive Chart Exploration	Strategy A	UAW over time		Drag gesture, temporal navigation
2	Search & Discover dApps	Strategy B	Inflow/Outflow		Search-based discovery, metric readability
3	Report-Based Overview	Strategy C	Transaction Fees, Gas Fees		Scroll-based scanning, metric findability

Table 5.1: Task scenario assignment across prototypes.

5.5 Measurement Instruments

Five parts (A through E) compose the questionnaire. Part A collects background data once. Parts B and C are repeated for each prototype. Parts D and E are collected once across all strategies. The complete item wording is reproduced in Appendix 8.

5.5.1 Task Performance and Experience Items

Each task scenario records the completion status on a three-level ordinal scale: completed independently, completed with help, or not completed. An overall prototype rating on a scale from 1 to 10 captures the participant’s summary assessment after each task.

After each task, participants respond to Likert items capturing perceived readability, data comprehension, and task-specific interaction quality. The items share a common core across tasks but are adapted to each task’s interaction pattern: Strategy A items reference chart interaction and trend identification, Strategy B items reference search result quality, and Strategy C items reference report scanning. Direct numerical comparison across strategies is limited to the readability item (E1), which uses identical wording in all three forms. Each task section concludes with an open-ended feedback prompt.

5.5.2 Modified System Usability Scale

Parts C.1, C.2, and C.3 each contain a five-item usability scale administered after the participant has reviewed the corresponding prototype. The scale draws on Brooke's [32] structure, retaining the alternating polarity pattern. Item wording has been adapted to reference visualization-specific aspects.

The original SUS comprises ten items. The instrument was reduced to five to manage questionnaire length in a within-subjects design where every participant rates all three prototypes. With task performance items, experience ratings, comparative assessments, and open-ended prompts, the full questionnaire already contained over sixty items. The five retained items cover perceived overview quality, organizational clarity, visual complexity, screen clutter, and user confidence. The reduction excludes learnability, consistency, and the need for prior knowledge; Section 7.4 evaluates the implications of this narrower scope. Scoring follows the standard SUS logic: positive items contribute (response - 1), negative items contribute (5 - response), and the sum is multiplied by 5 to yield a composite between 0 and 100. Because the instrument uses five rather than ten items, the resulting scores are not comparable to published SUS benchmarks but remain valid for within-study comparison. Section 7.4 addresses this constraint.

5.5.3 Comparative Rating and Final Assessment

Part D asks participants to rate all three prototypes side by side on two dimensions: data clarity and ease of use, each scored from 1 (poor) to 5 (excellent). The side-by-side format forces direct comparison under identical scales, providing the primary basis for cross-strategy comparison of information clarity and usability. Three context-specific preference items follow, each requiring participants to select the single best prototype for a given use case: getting a quick ecosystem overview, finding specific data, and recommending the application to a first-time blockchain user.

Part E concludes with an overall ranking of the three prototypes from best to least preferred, a free-text rationale, and a recommendation score on a scale from 1 to 10.

6 Results

Ten participants evaluated three visualization strategies through task scenarios, per-prototype usability scales, and cross-strategy comparative ratings. Strategy A received responses from all ten participants. Strategy B received ten submissions, though two are partially incomplete: one participant could not access the Figma prototype due to a broken link, and another completed the task but left the usability scale blank. Strategy C received ten submissions. Effective sample sizes vary by instrument and are noted where they differ from the full set.

6.1 Task Completion and Efficiency

Completion status was recorded on a three-level ordinal scale. An overall prototype rating on a scale from 1 to 10 accompanied each task. No time-on-task data was captured in the administered forms.

Strategy	n	Completed	With help	Not completed
A (Interactive Exploration)	10	10	0	0
B (Search & Discover)	10	8	1	1
C (Report Overview)	10	9	0	0

Table 6.1: Task completion status per strategy.

Strategy A achieved a 100 percent completion rate. Strategy B produced the only cases of incomplete or assisted execution: one participant could not access the prototype, a second needed assistance with metric switching, and a third submitted the task as completed but left all usability items blank. For Strategy C, nine of ten respondents completed the task independently; one left the completion field empty but provided experience ratings. The completion data separates into two patterns: tasks on the native application encountered no barriers, while tasks on Figma prototypes produced both technical failures and navigation difficulties. Whether this reflects properties of the visualization strategies or limitations of the prototyping medium cannot be determined from completion rates alone.

Strategy	n	Mean	Median	Min	Max
A	10	8.3	9.0	6	10
B	8	7.6	7.0	6	10
C	10	6.8	6.5	5	10

Table 6.2: Overall prototype ratings (1 to 10 scale).

Overall ratings reinforce the completion pattern. Strategy A leads at 8.3, with seven out of ten scores at 8 or above. Strategy B averages 7.6, with a compressed distribution between 6 and 8; its two highest scores (9 and 10) came from participants with expert-level blockchain knowledge. Strategy C received the lowest mean at 6.8, with a median of 6.5. Nine out of ten ratings fall between 5 and 8; one participant, the most experienced blockchain professional in the sample, assigned a 10, creating the widest spread in the dataset.




Figure 6.1: Distribution of overall prototype ratings per strategy.

Three themes appear across the open-ended responses. Strategy A drew praise for the stacked card layout and chart interactions but received criticism for the calendar-based time filter.

Strategy B's dominant theme was navigation difficulty: metric selection controls were hard to locate, and the Figma prototype's limited interactivity compounded this issue. Strategy C's information density drew polarized reactions: one expert praised the metric summaries, while several others described the layout as "dense," "heavy for mobile," or "overwhelming."

6.2 Usability Scale Comparison

The five-item modified SUS was administered once per prototype. Effective sample sizes are ten for Strategy A, eight for Strategy B, and ten for Strategy C.

Strategy	n	Mean	Median	SD	Min	Max
A	10	78.0	80.0	17.7	45	100
B	8	69.4	67.5	17.8	50	95
C	10	59.5	60.0	19.4	30	90

Table 6.3: Modified SUS composite scores per strategy.

The three strategies form a descending sequence with comparable within-group variance (SD 17.7 to 19.4). The gap between A and C spans roughly eighteen points. Strategy B occupies a middle position closer to A than to C. Breaking these composite scores into their constituent items reveals which aspects of usability drive the overall differences.

Item	Polarity	A (n=10)	B (n=8)	C (n=10)
Good ecosystem overview	Positive	4.10	4.00	3.80
Too complex / confusing	Negative	1.80	2.25	2.50
Well organized	Positive	4.30	3.88	3.50
Screens felt cluttered	Negative	2.00	2.12	2.70
I felt confident using it	Positive	4.00	3.38	2.80

Table 6.4: Item-level means across the five modified SUS items. Higher values indicate stronger agreement; for negative items, lower scores are preferable.



Figure 6.2: Mean ratings per modified SUS item across the three strategies.

Item-level analysis shows where the strategies diverge most sharply. The overview item produces modest differences (3.80 to 4.10): participants generally felt each strategy provided a reasonable ecosystem overview. Organization and confidence show larger gaps. Strategy A's organization rating of 4.30 drops to 3.50 for Strategy C. The confidence item shows the steepest decline: 4.00 for A to 2.80 for C. Perceived complexity and clutter track this pattern from the opposite direction, both rising from A to C.

SUS Scores by Expertise Subgroup

Splitting the sample into participants with advanced or expert blockchain knowledge (P1, P4, P8, P9, P10; $n = 5$) and those with intermediate or lower familiarity (P2, P3, P5, P6, P7; $n = 5$) quantifies how domain expertise moderates the usability gap across strategies.

Strategy	Expert/Advanced			Intermediate or lower		
	n	Mean	Range	n	Mean	Range
A	5	85.0	75–100	5	71.0	45–95
B	4	78.8	50–95	4	60.0	50–75
C	5	70.0	50–90	5	49.0	30–60

Table 6.5: Modified SUS composite scores by expertise subgroup.

The gap between subgroups widens as information density increases. On Strategy A, expert participants scored 14 points above non-experts. On Strategy C, the gap expands to 21 points. With five participants per subgroup, this pattern cannot be treated as a reliable estimate. It is consistent with the interpretation that domain knowledge partially compensates for the extraneous load that high-density layouts impose, but replication with a larger sample would be needed to assess whether the moderation effect holds beyond this dataset.

A participant-level pattern emerges across strategies. One participant with advanced blockchain knowledge rated all three prototypes consistently high (90, 95, 90). At the other end, the two lowest Strategy C scores (30 and 35) came from participants who had rated Strategy A at 70 and 90 respectively, neither with expert-level blockchain experience. The drop from A to C is steeper for participants outside the core expert audience than for those with regular exposure to blockchain analytics tools.

6.3 Information Clarity Across Strategies

Information clarity is captured through per-task experience items and the cross-prototype comparative rating from Part D. Only the readability item (E1) uses identical wording across all three forms.

Per-Task Experience Ratings

Strategy A produces the highest ratings across all items. Readability received near-unanimous agreement at 4.90. The lowest-rated item is data comprehension (E2) at 3.90, pulled down by one participant with no prior blockchain experience. Strategy B’s readability (4.56) approaches Strategy A’s level, but comprehension and findability both drop to 3.22, the

Strategy	Item	Mean	SD	n
A	E1: Data was readable on mobile	4.90	0.32	10
	E2: I could quickly understand the data	3.90	1.20	10
	E3: Easy to find what I needed	4.50	0.85	10
	E4: Information amount felt right	4.60	0.70	10
	E5: Charts helped me see trends	4.40	0.97	10
B	E1: Data was readable on mobile	4.56	0.73	9
	E2: I could quickly understand the data	3.22	1.30	9
	E3: Easy to find what I needed	3.22	0.97	9
	E4: Search results were helpful	3.89	0.60	9
C	E1: Data was readable on mobile	3.80	0.92	10
	E2: Report gave good overview	3.60	0.84	10
	E3: Easy to scan information	3.40	1.17	10
	E4: Found what I needed quickly	3.30	0.67	10

Table 6.6: Per-task experience item means (1 = Strongly Disagree, 5 = Strongly Agree).

widest gap between any readability score and its accompanying comprehension score. The search-specific item (E4: 3.89) scores higher, suggesting the search mechanism functioned adequately even when surrounding navigation did not. Strategy C records the lowest means on every item. The findability item (E4) at 3.30 with a standard deviation of 0.67 indicates moderate difficulty across participants, though with more spread than Strategy A or B on any single item.

Cross-Strategy Readability

The readability item (E1) declines from Strategy A (4.90) through B (4.56) to C (3.80), tracking the increase in information density. Five participants who rated Strategy A at 5 dropped to 3 on Strategy C. Three participants rated all three strategies at 5 or above, all with advanced or expert blockchain experience. The decline is concentrated among participants with intermediate or lower domain familiarity.

Comparative Ratings: Data Clarity and Ease of Use

Part D asked participants to rate all three prototypes on data clarity and ease of use using identical scales, which were collected after interaction with all three prototypes.

Strategy A leads on both dimensions. Its ease of use (4.80) is the highest mean in the dataset.

Participant	Crypto Knowledge	Strategy A	Strategy B	Strategy C
P1	Advanced	5	5	5
P2	Intermediate	4	–	4
P3	Basic	5	5	3
P4	Advanced	5	5	5
P5	Intermediate	5	5	3
P6	None	5	3	4
P7	Basic	5	4	3
P8	Expert	5	4	3
P9	Expert	5	5	3
P10	Advanced	5	5	5

Table 6.7: Cross-strategy readability (E1) by participant. Cells marked with – indicate missing data.

Dimension	Mean A	SD A	Mean B	SD B	Mean C	SD C
Data Clarity	4.50	0.71	3.67	0.50	3.30	0.82
Ease of Use	4.80	0.42	3.56	1.01	3.00	0.67
<i>n</i>	10		9		10	

Table 6.8: Part D comparative ratings (1 = Poor, 5 = Excellent).

Data clarity at 4.50 falls slightly below, producing a pattern where the interface is perceived as easier to operate than to extract information from. Strategy B reverses this narrowly: clarity (3.67) sits above ease (3.56). Strategy C shows the widest gap: data clarity at 3.30 exceeds ease of use at 3.00, the largest clarity-over-ease margin across all three strategies.




Figure 6.3: Comparative ratings for data clarity and ease of use per strategy. Strategy C shows the largest divergence between the two dimensions.

The direction of the gap shifts across strategies. Strategy A is rated easier to use (4.80) than clear (4.50). Strategy C is rated clearer (3.30) than easy to use (3.00). This reversal means that progressive disclosure favors operational fluency, while the high-density report layout preserves informational content at the cost of navigational effort. Strategy B occupies a transitional position where clarity and ease are approximately balanced.

6.4 Comparative Ranking and Preference

Ten participants completed the preference and ranking sections.

Context-Specific Preference

Strategy A dominates all three categories. Strategy B received only two votes for finding specific data, reflecting its search-oriented design. No participant selected Strategy B for quick overview or beginner suitability, meaning its narrative scaffolding did not translate into perceived accessibility.

Use case	Strategy A	Strategy B	Strategy C
Quick ecosystem overview	8	0	2
Finding specific data	7	2	1
Recommending to a beginner	9	0	1

Table 6.9: Context-specific preference votes (n = 10).

Overall Ranking

Rank	Strategy A	Strategy B	Strategy C
1st (Best)	8	0	2
2nd	2	6	2
3rd (Least preferred)	0	4	6

Table 6.10: Overall ranking distribution (n = 10).

Eight participants ranked Strategy A first; none placed it last. Strategy B settled into second place for six respondents. Strategy C was ranked last by six and first by two. Both first-place votes for Strategy C carry caveats: one participant cited market entry potential rather than usability or clarity as ranking criteria, while the other could not access Strategy B and ranked it last by default.

Recommendation Score

Metric	Value
Mean	8.80
Median	9.00
Min / Max	6 / 10
Promoters (9–10)	5
Passives (7–8)	4
Detractors (1–6)	1
NPS	+40

Table 6.11: Recommendation score distribution (1 to 10 scale, n = 10).

The NPS of +40 reflects positive sentiment. The single detractor (score of 6) is the participant who experienced the Figma access issue. At n = 10, the NPS value is sensitive to individual responses and not interpretable as a stable estimate; it is reported here for completeness rather than as a substantive finding.




Figure 6.4: Overall ranking distribution across the three strategies ($n = 10$).

Qualitative Themes

Open-ended responses surface three recurring themes. The first is the role of information hierarchy: participants who ranked Strategy A first pointed to the readability of its home screen and the logical progression from summary to detail. The second is the tension between density and comprehension: Strategy C was described as "powerful" by experts and "overwhelming" by others. The third is the confounding effect of prototype fidelity: multiple participants noted that Strategy A felt "finished" or "well thought out," language pointing to implementation quality rather than visualization design. These remarks indicate that some portion of Strategy A's ranking advantage reflects implementation medium rather than strategy-level choices.

7 Discussion

The results present a consistent ranking across instruments: Strategy A scores highest, Strategy B occupies a middle position, and Strategy C trails on most measures. That ranking alone would be a thin answer to the research question. The interpretation that follows argues that the pattern is more layered than the ranking suggests, that usability and information clarity do not move in lockstep, and that the ranking itself is shaped by a confound the evaluation can identify but not fully resolve.

7.1 Interpretation of Comparative Findings

The research question asks how different visualization strategies for blockchain network data perform on mobile devices in terms of usability and information clarity. Two instruments provide the most direct answer: the modified SUS composite, which captures perceived usability, and the Part D comparative rating, which separates data clarity from ease of use under identical scales.

On usability, the 18.5-point gap between A and C approaches the within-strategy standard deviation, while the 8.6-point gap between A and B is more modest. Confidence shows the steepest decline across strategies (4.00 to 2.80), in line with the cognitive load trajectory predicted in Chapter 3: as simultaneously visible information elements increase, working memory demands rise and self-reported confidence drops [6].

On information clarity, the Part D rating complicates the picture. Clarity gaps between strategies (1.20 points from A to C) are smaller than ease of use gaps (1.80 points). Strategy C preserves more of its standing when participants assess the data itself rather than the experience of accessing it. A findability item at 3.30 captures this: locating specific data required effort; yet the same participants gave Strategy C moderate marks on overview quality (3.60) and readability (3.80). Present and structured, but harder to reach.

More informative than the ranking itself is a reversal in the relationship between clarity and ease. Strategy A is rated easier to use (4.80) than clear (4.50). Strategy C is rated clearer (3.30) than easy to use (3.00). Strategy B sits at approximately the same level. Rather than

occupying different positions along a single quality axis, the three strategies distribute their strengths differently across the two constructs that the research question targets.

Strategy A's progressive disclosure defers detail behind card taps and bottom sheets, making the interface operationally smooth. A user who navigates the home screen and scans the KPI values has completed a usable interaction without necessarily encountering the stacked area charts or ranked bar views that contain the analytical substance. The home screen's KPI cards present numbers without visual context: a transaction count without a trend line, a wallet figure without a distribution chart. Progressive disclosure, in other words, tends to optimize the operational layer at the expense of the analytical one.

Strategy C reverses this dynamic. Charts, KPI grids, and wallet visualizations appear on the initial screen simultaneously. Participants registered this informational presence in the clarity ratings. But the visual density also generates the clutter reflected in the usability items. Cognitive load theory accounts for this with precision: the intrinsic load of the blockchain data remains constant across strategies, while the extraneous load imposed by the presentation changes. Strategy C's simultaneous display produces extraneous load that competes with the analytical task.

Strategy B's intermediate position is not simply an average of the extremes. The narrative text was designed to reduce interpretive effort, but most participants described the prototype in terms of navigation difficulty rather than narrative quality. The search function worked (E4: 3.89), but metric switching controls and overall flow produced confusion that overshadowed whatever clarity the narrative text contributed. The theoretical advantage of offloading cognitive effort through text-based scaffolding did not manifest because the interaction layer introduced enough friction to neutralize it.

Experience level cuts across all three strategies but is most pronounced in Strategy C. SUS gaps between expert and non-expert participants reach 21 points in Strategy C (Table 6.5), compared to 14 points in Strategy A. The data is consistent with domain knowledge functioning as a buffer against extraneous load: a user who already knows what UAW means would need fewer cognitive resources to interpret the charts, leaving more capacity for navigation. Strategy A's lower information density reduces this effect because fewer elements compete for attention regardless of the user's background.

7.2 Trade-offs Between the Strategies

Each strategy rests on a specific assumption about how blockchain data should be presented on a mobile screen. The evaluation data allows these hypotheses to be assessed against observed outcomes.

Strategy A: Disclosure Depth versus Surface Comprehension

Strategy A's card-based progressive disclosure produced the highest usability scores on every instrument. The cost is visible in a specific pattern: ease of use exceeds clarity by a third of a scale point (4.80 vs. 4.50), a gap no other strategy shows.

The explanation connects to what progressive disclosure hides. Cleveland and McGill's [1] perceptual hierarchy favors the position-based encodings that Strategy A uses in its detail views, but those encodings are available only after a tap. The home screen presents numbers without visual context. For the quick health check that foundation teams perform between meetings, this may suffice. For analytical comparison across categories, the user must know that deeper views exist and actively seek them out. Progressive disclosure keeps the screen clean at the price of hiding connections between data layers. On desktop, adjacent panels make those connections spatially explicit. On mobile, they must be reconstructed mentally across sequential views, a task requiring productive schema construction: effortful, but directed toward comprehension rather than compensating for a poor interface, provided the user has the domain knowledge to connect the pieces.

Strategy B: Narrative Scaffolding versus Navigation Friction

Strategy B tested whether embedding interpretive text alongside charts would reduce cognitive effort. The evaluation does not confirm this benefit. Strategy B's clarity rating of 3.67 matches its overall middle position, and its comprehension item (3.22) falls substantially below Strategy A's.

Two factors neutralized the narrative's potential. The chip-based filter system occupies the top of the screen, a placement Snyder et al. [26] identify as problematic for thumb-operated interfaces. Several participants could not locate the filter controls without assistance. That alone caps the clarity gains the narrative text might otherwise deliver: a navigation barrier encountered before the content is reached nullifies whatever interpretive scaffolding follows.

The Figma prototype compounds the problem further, since filter state changes produce discrete screen jumps rather than smooth transitions, and the resulting disorientation sits on top of the strategy-inherent navigation issues. The collected data cannot separate the two sources of friction.

Strategy B's search function represents a partial success. The search item (E4: 3.89) scored highest among Strategy B's experience items, and the two preference votes it received for finding data reflect this. The strategy's contribution may lie less in narrative scaffolding than in its search-first architecture: when users could bypass the navigation layer and reach content directly, the prototype performed closer to its design intent.

Strategy C: Informational Presence versus Operational Cost

Strategy C inverts Strategy A's approach, placing charts, KPI grids, and wallet visualizations on screen simultaneously. The evaluation confirms one half of the design hypothesis and rejects the other. The clarity rating (3.30) does not collapse as far as the ease of use rating (3.00), producing the only strategy where perceived clarity exceeds perceived ease. The problem is not informational absence but operational access.

The clutter item (2.70) identifies where the cost originates. Horak et al. [23] propose responsive levels of detail as a solution for constrained screen space. Strategy C does not implement this adaptation; each element renders at full complexity regardless of how many others share the screen, exceeding the clutter threshold Sarikaya et al. [13] identify. The chart carousel introduces a different tradeoff. Brehmer et al. [24] found that spatial juxtaposition outperforms sequential presentation on mobile. The carousel is sequential by design: comparing temporal trends against part-to-whole proportions requires holding the first chart in memory while swiping to the second. One participant confirmed this: "similar data blocks on the sibling screen have different representation, little confusing when swiping."

The available data suggests the expertise moderation effect is strongest here. The SUS gap between expert and non-expert participants reaches 21 points on Strategy C, compared to 14 on Strategy A (Table 6.5). In this sample, dense displays appear to reward prior knowledge and penalize its absence. The design serves a narrower audience, but within that audience, it delivers what Strategy A does not: immediate access to analytical depth without navigational intermediaries.

The Density Gradient

Across the three strategies, information density increases from A through B to C. Usability declines, clarity declines less steeply, and the clarity-ease relationship shifts from ease-over-clarity through parity to clarity-over-ease. The gradient suggests that the tradeoff is continuous rather than binary. A quick ecosystem health check benefits from Strategy A's restraint. A detailed category comparison may require the density that Strategy C provides, at least for users equipped to handle it.

7.3 Implications for Mobile Data Visualization

Several findings generalize beyond blockchain ecosystems. The data characteristics that make blockchain visualization difficult on mobile, temporal series, categorical breakdowns, hierarchical aggregation, and multiple concurrent metrics, recur in any domain where multidimensional performance data must be monitored on a phone.

Progressive Disclosure as Default

Strategy A's consistent advantage supports a practical recommendation: on mobile devices, progressive disclosure should serve as the default information architecture for complex data domains. The principle itself is not new [15]. What the evaluation adds is evidence about the granularity at which it operates most effectively on a phone: compressing the overview to a single chart and four numeric values, while placing all metric-specific visualizations behind explicit card taps.

The tradeoff documented in Section 7.2 suggests a refinement. Pure numeric KPI cards without embedded trend indicators may be too reductive for analytical tasks. Adding sparklines or directional arrows, following the responsive levels of detail proposed by Horak et al. [23] could increase clarity at the overview level without breaking the progressive disclosure structure.

Narrative Scaffolding and Spatial Competition

On a mobile screen, text competes with visualization for the same scarce spatial resource. A narrative paragraph displaces a chart, a legend, or whitespace that might have contributed

more to comprehension. Short, inline annotations attached to specific chart features occupy less space and provide context exactly where attention is directed. Munzner’s [21] effectiveness principle applies: the most valuable information should occupy the most perceptually salient channel. On a six-inch screen, spatial position is that channel, and dedicating it to running text may be a misallocation.

Information Density as Expertise-Dependent Variable

Strategy C’s dense layout produced SUS scores ranging from 30 to 90, a spread that correlates with blockchain knowledge more closely than any other participant characteristic. Information density on mobile dashboards should not be treated as a fixed design parameter. Adaptive approaches offer a path forward: an interface that opens at low density and expands on demand serves the broader audience, while an option to switch into a dense monitoring mode serves expert users who value simultaneous data access. Hoffswell et al. [22] documented responsive transformations that adapt visualizations to screen size; the same logic could apply to user expertise.

Fidelity Asymmetry as Methodological Evidence

Mixed-fidelity evaluation is common in industry contexts where resource constraints prevent the equal implementation of all design alternatives. The present study did not set out to measure the fidelity effect, and the data cannot decompose it from genuine design differences. Strategy A leads Strategy B by 8.6 SUS points and Strategy C by 18.5 points; what share of each gap reflects visualization design rather than implementation completeness remains unknown. Comparative visualization research on mobile devices rarely reports fidelity as a controlled variable. By holding the data domain constant and varying both visualization design and implementation fidelity simultaneously, this thesis exposes a confound that other studies may carry silently.

Transferable Findings

The design tensions identified in Section 3.2 recur in financial analytics, IoT monitoring, and AdTech dashboards. The clarity-ease reversal provides a transferable diagnostic: when a mobile dashboard scores higher regarding clarity than ease, the architecture is likely too

dense; when ease exceeds clarity, it may be hiding needed information. Measuring both dimensions separately is methodologically worthwhile for any comparative evaluation of mobile data interfaces.

This thesis contributes to the scarcity of mobile visualization research that Blumenstein et al. [5] identified, in three respects: a within-subjects comparison of structurally different strategies on the same data, the operationalization of information clarity as a dimension separate from usability, and documentation of the moderating role of domain expertise. The sample is small and the domain narrow, but the evaluation framework and the clarity-ease separation apply wherever mobile dashboards present high-dimensional data to users of varying expertise.

7.4 Limitations and Threats to Validity

Internal Validity

The most consequential threat is the fidelity asymmetry between the prototypes. Strategy A was implemented as a native iOS application with dynamic data, continuous gesture interaction, and haptic feedback. Strategies B and C were Figma prototypes with static screen variants. Multiple participants explicitly described Strategy A as "finished" or "most complete." Sonderegger and Sauer [29] found that visual aesthetics and perceived completeness influence usability ratings independently of interaction quality, suggesting the fidelity gap inflates Strategy A's scores by an unknown margin. The Part D ratings partially mitigate this through forced comparison under identical scales, but the underlying experience was shaped by the fidelity difference.

The asymmetry also yields its own findings. The consistent gap between Strategy A and the Figma prototypes across instruments that measure different constructs, usability, clarity, and preference, documents the degree to which implementation fidelity co-determines evaluation outcomes in comparative prototype research. Future studies that compare visualization strategies at mixed fidelity levels can use this documented margin as a calibration reference.

Each prototype received one task targeting its defining interaction paradigm. Tasks were not crossed with strategies. A participant's experience ratings reflect one scenario, not the full range of a prototype's capabilities. The modified SUS and Part D instruments compensate partially because they were administered after overall prototype exposure, not tied to a

single task.

The remote, unmoderated setting introduced variability: one participant tested Strategy A on a Mac rather than a phone, another could not access Strategy B due to a broken link. No counterbalancing was applied for form order, and submission timestamps suggest most participants completed in the sequence A, B, C. Any order effect would systematically favor whichever prototype was encountered first or last.

External Validity

Ten participants were purposively recruited from FiDi Tech and its professional network. The sample provides ecologically valid feedback from the target user population but limits generalizability. Blumenstein et al. [5] report a median sample size of twelve across 21 mobile visualization studies, placing this count near the lower range. Company affiliation introduces potential courtesy bias, particularly for Strategy A as the most advanced product artifact.

All prototypes used mock data, structurally realistic but not drawn from live blockchain activity. Real data exhibits volatility spikes, missing periods, and anomalous values that mock data smooths over. The restriction to a single ecosystem, Moonbeam within Polkadot, further limits the scope. Other architectures with different data structures or higher cardinality might expose different tradeoffs.

Construct Validity

The modified SUS uses five items rather than the standard ten. The retained items cover overview quality, complexity, organization, clutter, and confidence; items addressing learnability, consistency, and the need for prior knowledge are absent. Scores are valid for within-study comparison but not for comparison to published benchmarks calibrated on the ten-item formulation.

Information clarity combines per-task experience ratings with the Part D data clarity item. The per-task items differ across strategies, limiting direct comparison to the single readability item. The Part D clarity rating collapses the multi-level construct described by Borgo et al. [28], elementary, intermediate, and comprehensive reading, into a single scale point. A more granular assessment with separate items per reading level would strengthen construct

validity.

The absence of objective task accuracy data further constrains measurement. The evaluation captured subjective ratings and completion status but did not ask participants to report specific data values extracted from the prototypes. Whether a participant who rated clarity at 4 actually read the chart correctly is unknown. Adding value-extraction tasks with verifiable answers would ground the clarity construct in behavioral evidence rather than in self-report alone.

8 Conclusion and Outlook

This thesis asked how different visualization strategies for blockchain network data perform on mobile devices in terms of usability and information clarity. Three strategies were designed, prototyped, and evaluated by ten participants in a within-subjects comparative study. Progressive disclosure (Strategy A) produced the strongest results across all instruments. Narrative scaffolding with search-based access (Strategy B) occupied a middle position, and the high-density report layout with carousel navigation (Strategy C) scored lowest on usability while retaining comparatively more of its clarity rating. The answer to the research question, however, is not this ranking but the structural relationship beneath it: as information density increases from A through B to C, ease of use drops faster than data clarity. The two dimensions do not collapse onto a single scale but define a tradeoff space whose optimal position depends on the user's expertise and analytical intent.

The work contributes on three levels. Most concretely, it provides three structurally distinct visualization strategies for the same mobile blockchain data domain, filling a design space where no comparable alternatives existed in published research. The evaluation framework is a second contribution: measuring usability and information clarity as separate constructs proved necessary because a unified quality metric would have masked the clarity-ease reversal that turned out to be the most informative finding. Empirically, the study adds evidence on mobile visualization tradeoffs drawn from a real industry context, addressing the scarcity of applied mobile visualization research that Blumenstein et al. [5] identified. The boundary conditions under which these contributions hold, including fidelity asymmetry, sample size, and construct coverage, are documented in Section 7.4.

Outlook

Future work should address the fidelity confound first. Implementing all three strategies as native applications with equivalent interaction quality, dynamic data, and gesture support would isolate visualization design effects from implementation medium effects. The evaluation data suggests this is worth the investment: qualitative feedback repeatedly attributes

part of Strategy A's perceived advantage to its native feel rather than its information architecture.

A larger and more diverse sample would strengthen the findings on expertise moderation. The current data indicates that domain knowledge buffers against the usability costs of high-density displays, but ten participants across three experience tiers cannot establish the shape of that relationship. Recruiting participants outside the blockchain domain, from adjacent fields such as financial analytics or IoT monitoring, would test whether the observed tradeoffs generalize to other high-dimensional mobile data contexts.

Strategy A is scheduled to be connected to FiDi Tech's production API, replacing the mock data module with live blockchain feeds. This transition will introduce the volatility, missing values, and anomalous patterns that production dashboards must handle. Whether the progressive disclosure approach remains effective when KPI values spike unpredictably, or whether the dense layout's simultaneous display becomes more valuable when users need to correlate rapid changes across metrics are questions that static mock data cannot answer.

Longitudinal evaluation would address a gap that the current study design leaves open. A single-session assessment captures first impressions and initial learnability but not the learning curves that develop over repeated use. Strategy C's dense layout may become more navigable as users internalize its structure, narrowing the usability gap with Strategy A over time. Conversely, Strategy A's progressive disclosure may frustrate experienced users who begin to perceive the tapping as overhead rather than guidance. Whether the strategies' relative positions shift with familiarity is an empirical question that requires weeks rather than minutes of observation.

The adaptive density approach suggested in Section 7.3 deserves concrete investigation. An interface that opens at Strategy A's level of restraint and allows users to switch into a dense monitoring mode resembling Strategy C's layout would test whether combining the two endpoints of the density gradient produces better outcomes than either alone. One participant independently suggested exactly this adaptation during the evaluation. Testing such a hybrid would move the findings closer to actionable design guidance.

Beyond density adaptation, the clarity-ease reversal itself suggests a diagnostic tool for mobile dashboard evaluation. Measuring both dimensions separately and comparing their relationship could serve as a lightweight quality indicator during iterative design: when ease substantially exceeds clarity, the interface may be hiding needed information; when clarity exceeds ease, the interface may be overwhelming its users. Applying this diagnostic across

domains and screen sizes would test whether the reversal pattern is specific to blockchain data or reflects a general property of information-dense mobile interfaces.

The research question framed the problem as a comparison. The answer reframes it as a tradeoff. Mobile blockchain visualization does not have a single best strategy. It has a design space defined by the tension between operational ease and informational depth, moderated by the expertise of the person holding the phone.

Bibliography

- [1] W. S. Cleveland and R. McGill, “Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods,” *Journal of the American Statistical Association*, vol. 79, no. 387, pp. 531–554, 1984.
- [2] F. M. Garrels, “Projektbericht: FiDi Mobile Analytics,” Hochschule Osnabrück, Osnabrück, Projektbericht 1, 2025.
- [3] N. Tovanich, N. Heulot, J.-D. Fekete, and P. Isenberg, “Visualization of Blockchain Data: A Systematic Review,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 7, pp. 3135–3152, Jul. 2021. [Online]. Available: <https://ieeexplore.ieee.org/document/8945380/>
- [4] N. Tovanich, N. Heulot, J.-D. Fekete, and P. Isenberg, “A Systematic Review of Online Bitcoin Visualizations,” in *Eurographics Conference on Visualization (EuroVis), Short Papers*, 2019.
- [5] K. Blumenstein, C. Niederer, M. Wagner, G. Schmiedl, A. Rind, and W. Aigner, “Evaluating Information Visualization on Mobile Devices: Gaps and Challenges in the Empirical Evaluation Design Space,” in *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization (BELIV '16)*. ACM, 2016, pp. 125–132.
- [6] J. Sweller, “Cognitive Load During Problem Solving: Effects on Learning,” *Cognitive Science*, vol. 12, no. 2, pp. 257–285, Apr. 1988. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1207/s15516709cog1202_4
- [7] E. A. Napoli, I. Spada, N. Romani, V. Gatteschi, and C. Schifanella, “Smart Contract Visualization: Solutions and Challenges,” in *Proceedings of the Seventh Distributed Ledger Technology Workshop (DLT 2025)*, vol. 4105. CEUR-WS.org, 2025. [Online]. Available: <https://ceur-ws.org/Vol-4105/paper10.pdf>
- [8] International Organization for Standardization, “ISO 9241-11:2018 Ergonomics of Human-System Interaction – Part 11: Usability: Definitions and Concepts,” 2018, published: International Standard. [Online]. Available: <https://www.iso.org/standard/>

63500.html

- [9] S. Ferretti and G. D’Angelo, “On the Ethereum Blockchain Structure: A Complex Networks Theory Perspective,” *Concurrency and Computation: Practice and Experience*, vol. 32, no. 12, 2020.
- [10] G. Wood, “Polkadot: Vision for a Heterogeneous Multi-Chain Framework,” Web3 Foundation, Whitepaper, 2016. [Online]. Available: <https://polkadot.com/papers/Polkadot-whitepaper.pdf>
- [11] J. Burdges, A. Cevallos, P. Czaban, R. Habermeier, S. Hosseini, F. Lama, H. K. Alper, X. Luo, F. Shirazi, A. Stewart, and G. Wood, “Overview of Polkadot and its Design Considerations,” May 2020, arXiv:2005.13456 [cs]. [Online]. Available: <http://arxiv.org/abs/2005.13456>
- [12] Web3 Foundation, “Web3 and Polkadot,” 2025, published: Polkadot Wiki. [Online]. Available: <https://wiki.polkadot.com/general/web3-and-polkadot/>
- [13] A. Sarikaya, M. Correll, L. Bartram, M. Tory, and D. Fisher, “What Do We Talk About When We Talk About Dashboards?” *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 682–692, Jan. 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8443395/>
- [14] Y. Riady, “How to Calculate DAU/WAU/MAU in Crypto and DeFi,” 2025, published: Formo Blog. [Online]. Available: <https://formo.so/blog/web3-active-users-dau-wau-mau>
- [15] B. Shneiderman, “The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations,” in *Proceedings of the IEEE Symposium on Visual Languages*. IEEE, 1996, pp. 336–343.
- [16] X. Yue, X. Shu, X. Zhu, X. Du, Z. Yu, D. Papadopoulos, and S. Liu, “BitExTract: Interactive Visualization for Extracting Bitcoin Exchange Intelligence,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 162–171, 2019.
- [17] B. Putz, F. Böhm, and G. Pernul, “HyperSec: Visual Analytics for blockchain security monitoring,” Mar. 2021, arXiv:2103.14414 [cs]. [Online]. Available: <http://arxiv.org/abs/2103.14414>
- [18] M. Sipos and S. Szénási, “Lightweight Blockchain Simulation with Transaction Graph Visualizer,” *Innovation in Science and Technology*, vol. 3, no. 1, 2024.

-
- [19] S. Burigat and L. Chittaro, “On the effectiveness of Overview+Detail visualization on mobile devices,” *Personal and Ubiquitous Computing*, vol. 17, no. 2, pp. 371–385, Feb. 2013. [Online]. Available: <http://link.springer.com/10.1007/s00779-011-0500-3>
- [20] J. Heer and M. Bostock, “Crowdsourcing Graphical Perception: Using Mechanical Turk to Assess Visualization Design,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, 2010, pp. 203–212.
- [21] T. Munzner, *Visualization Analysis and Design*, 0th ed. A K Peters/CRC Press, Dec. 2014. [Online]. Available: <https://www.taylorfrancis.com/books/9781466508934>
- [22] J. Hoffswell, W. Li, and Z. Liu, “Techniques for Flexible Responsive Visualization Design,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. ACM, 2020, pp. 1–13.
- [23] T. Horak, W. Aigner, M. Brehmer, A. Joshi, and C. Tominski, “Responsive Visualization Design for Mobile Devices,” in *Mobile Data Visualization*, 1st ed. Boca Raton: Chapman and Hall/CRC, Nov. 2021, pp. 33–66. [Online]. Available: <https://www.taylorfrancis.com/books/9781003090823/chapters/10.1201/9781003090823-2>
- [24] M. Brehmer, B. Lee, P. Isenberg, and E. K. Choe, “A Comparative Evaluation of Animation and Small Multiples for Trend Visualization on Mobile Phones,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 1, pp. 364–374, 2020.
- [25] M. Brehmer, B. Lee, J. Stasko, and C. Tominski, “Interacting with Visualization on Mobile Devices,” in *Mobile Data Visualization*, 1st ed., B. Lee, R. Dachsel, P. Isenberg, and E. K. Choe, Eds. Boca Raton: Chapman and Hall/CRC, Nov. 2021, pp. 67–110. [Online]. Available: <https://www.taylorfrancis.com/books/9781003090823/chapters/10.1201/9781003090823-3>
- [26] L. S. Snyder, R. A. Rossi, E. Koh, J. Heer, and J. Hoffswell, “Interaction Techniques for Exploratory Data Visualization on Mobile Devices,” Apr. 2024, arXiv:2404.11602 [cs]. [Online]. Available: <http://arxiv.org/abs/2404.11602>
- [27] M. Kurosu and A. Hashizume, “Usability and User Experience (UX) of Human-Automation Interaction: Concept and Measurement Method,” in *Human-Automation Interaction: Manufacturing, Services and User Experience*, ser. Automation, Collaboration, & E-Services, V. G. Duffy, M. Lehto, Y. Yih, and R. W. Proctor, Eds. Springer, 2023, vol. 10.
- [28] R. Borgo, B. Lee, B. Bach, S. Fabrikant, R. Jianu, A. Kerren, S. Kobourov, F. McGee,
-

- L. Micalef, T. Von Landesberger, K. Ballweg, S. Diehl, P. Simonetto, and M. Zhou, “Crowdsourcing for Information Visualization: Promises and Pitfalls,” in *Evaluation in the Crowd. Crowdsourcing and Human-Centered Experiments*. Springer, 2017.
- [29] A. Sonderegger and J. Sauer, “The Influence of Prototype Fidelity and Aesthetics of Design in Usability Tests: Effects on User Behaviour, Subjective Evaluation and Emotion,” *Applied Ergonomics*, vol. 41, no. 1, pp. 130–140, 2010.
- [30] J. Lazar, J. H. Feng, and H. Hochheiser, *Research Methods in Human-Computer Interaction*, 2nd ed. Morgan Kaufmann, 2017.
- [31] J. Blattgerste and J. Behrends, “A Web-Based Analysis Toolkit for the System Usability Scale,” in *Proceedings of the 15th International Conference on PErvasive Technologies Related to Assistive Environments (PETRA ’22)*. ACM, 2022, pp. 237–246.
- [32] J. Brooke, “SUS: A Quick and Dirty Usability Scale,” in *Usability Evaluation in Industry*, P. W. Jordan, B. Thomas, I. L. McClelland, and B. Weerdmeester, Eds. Taylor & Francis, 1996, pp. 189–194.

Appendix

A Task Scenario Descriptions

The following task descriptions were provided to participants as part of the Google Forms questionnaire. Each task specifies a sequence of interaction steps, a target metric, and the prototype on which the task was to be performed.

Task 1: Interactive Chart Exploration (Strategy A)

1. From the home screen, locate and tap on the search functionality.
2. Type “Beamswap” into the search field.
3. Select “Beamswap” from the search results.
4. On the detail page, find the UAW (Unique Active Wallets) chart.
5. Locate the time range selector and change it to “1 Year.”
6. Interact with the chart to explore UAW trends over time.

Target metric: UAW over time.

Task 2: Search and Discover dApps (Strategy B)

1. From the home screen, tap on the “Categories” tab in the bottom navigation.
2. Locate and tap the search icon at the top of the screen.
3. Type the letter “O” into the search field and wait for results.
4. From the search results, find and select “Outmine.”
5. On the Outmine detail page, locate the Inflow/Outflow visualization.
6. Observe how the token movement data is displayed.

Target metric: Inflow/Outflow.

Task 3: Report-Based Overview (Strategy C)

1. From the home screen, tap on the “Categories” tab in the bottom navigation.
2. Browse the available categories and tap on “Gaming.”
3. Scroll through the report-style layout with its sections.
4. Observe how different metrics are organized in the report view.
5. Locate the section showing Transaction Fees data.
6. Note how Gas Fees are visualized alongside other metrics.

Target metrics: Transaction Fees, Gas Fees.

B Questionnaire

The complete questionnaire comprised five parts (A through E), administered across three Google Forms. Part A was collected once. Parts B and C were repeated per prototype. Parts D and E were collected once across all prototypes.

Part A: Participant Background

Item	Response format
Date	Open
Job title	Open
Age	Open
Occupation	Open
<i>Experience with Data Visualization Apps</i>	
How often do you use apps or dashboards with charts and data visualizations?	Never / Rarely / Monthly / Weekly / Daily
<i>Blockchain/Crypto Experience</i>	
How would you rate your blockchain and crypto knowledge?	None / Basic / Intermediate / Advanced / Expert
Have you used blockchain analytics tools such as DappRadar, DefiLlama, or Etherscan?	Yes / No (if yes: specify)
<i>Mobile App Usage</i>	
How often do you use mobile apps for financial or analytics data?	Never / Rarely / Monthly / Weekly / Daily

Table 8.1: Part A: Participant background items.

Part B: Task Scenarios with Experience Items

After each task scenario (see Appendix 8), participants responded to a set of Likert-scaled experience items. Task completion status and an overall prototype rating (1 to 10) were also recorded per task. The experience items share a common structure across tasks but are adapted to each task's interaction pattern. Table 8.2 through Table 8.4 list the exact items administered per task.

ID	Item	Scale
<i>Task performance</i>		
	Completed?	Yes / With help / No
<i>Experience items (1 = Strongly Disagree, 5 = Strongly Agree)</i>		
E1	Data was readable on mobile.	5-point Likert
E2	I could quickly understand the data.	5-point Likert
E3	Easy to find what I needed.	5-point Likert
E4	Information amount felt right.	5-point Likert
E5	Charts helped me see trends.	5-point Likert

Table 8.2: Part B, Task 1 (Strategy A): Experience items for interactive chart exploration.

ID	Item	Scale
<i>Task performance</i>		
	Completed?	Yes / With help / No
<i>Experience items (1 = Strongly Disagree, 5 = Strongly Agree)</i>		
E1	Data was readable on mobile.	5-point Likert
E2	I could quickly understand the data.	5-point Likert
E3	Easy to find what I needed.	5-point Likert
E4	Search results were helpful.	5-point Likert

Table 8.3: Part B, Task 2 (Strategy B): Experience items for search-based discovery.

ID	Item	Scale
<i>Task performance</i>		
	Completed?	Yes / With help / No
<i>Experience items (1 = Strongly Disagree, 5 = Strongly Agree)</i>		
E1	Data was readable on mobile.	5-point Likert
E2	Report gave good overview.	5-point Likert
E3	Easy to scan information.	5-point Likert
E4	Found what I needed quickly.	5-point Likert

Table 8.4: Part B, Task 3 (Strategy C): Experience items for report-based scanning.

Each task section concluded with an open-ended prompt: “What worked well / what was difficult with this visualization approach?”

Part C: Modified System Usability Scale

The following five items were administered once per prototype (Parts C.1, C.2, and C.3). Response scale: 1 = Strongly Disagree to 5 = Strongly Agree. The instrument retains the alternating polarity pattern of Brooke’s original SUS [32] but reduces the item count from ten to five to manage questionnaire length in a within-subjects design with three prototypes. Items cover perceived overview quality, complexity, organization, visual clutter, and user confidence.

No.	Item	Polarity
1	The visualizations gave me a good ecosystem overview.	Positive
2	The information architecture was unnecessarily complex.	Negative
3	Data was organized in a logical, explorable way.	Positive
4	The screens felt cluttered or overwhelming.	Negative
5	I felt confident exploring the blockchain ecosystem.	Positive

Table 8.5: Part C: Modified SUS items (5-item variant). Scoring: positive items contribute (response – 1), negative items contribute (5 – response). The sum is multiplied by 5 to yield a composite score between 0 and 100.

Each Part C section concluded with an open-ended prompt: “Overall impression of this

visualization approach.”

Part D: Comparative Analysis

Dimension	Response format
Data Clarity	1 (Poor) to 5 (Excellent), per prototype
Ease of Use	1 (Poor) to 5 (Excellent), per prototype

Table 8.6: Part D: Cross-prototype comparative rating dimensions.

Best prototype for:	Response format
Getting a quick ecosystem overview	P1 / P2 / P3
Finding specific data	P1 / P2 / P3
Recommending to a first-time blockchain user	P1 / P2 / P3

Table 8.7: Part D: Context-specific preference items.

Part E: Final Assessment

Item	Response format
Overall Ranking	Rank 1 (best) to 3 (least preferred) per prototype
Why did you rank them this way?	Open-ended
Any suggestions to improve blockchain visualization on mobile?	Open-ended
Recommendation score	1 to 10

Table 8.8: Part E: Final assessment and ranking items.

C Raw Evaluation Data

The following tables present the complete per-participant response data collected across the three evaluation forms. Cells marked with – indicate missing responses. Participant

identifiers (P1–P10) are consistent across all tables.

Participant Demographics

ID	Role	Crypto Knowledge
P1	CEO	Advanced
P2	Data Engineer	Intermediate
P3	Student	Basic
P4	Developer	Advanced
P5	Data Analyst	Intermediate
P6	Student	None
P7	Techfounder	Basic
P8	CEO	Expert
P9	Founder	Expert
P10	CTO	Advanced

Table 8.9: Participant demographics.

Task Completion and Overall Ratings

ID	Task Completed			Overall Rating		
	A	B	C	A	B	C
P1	Yes	Yes	Yes	8	–	8
P2	Yes	No	Yes	6	–	6
P3	Yes	Yes	Yes	9	7	6
P4	Yes	Yes	Yes	7	8	5
P5	Yes	Yes	Yes	9	7	7
P6	Yes	Yes	Yes	7	6	7
P7	Yes	Yes	–	9	7	6
P8	Yes	With help	Yes	9	7	6
P9	Yes	Yes	Yes	9	9	7
P10	Yes	Yes	Yes	10	10	10

Table 8.10: Task completion status and overall prototype ratings (1–10) per participant.

Experience Items

ID	Strategy A					Strategy B				Strategy C			
	E1	E2	E3	E4	E5	E1	E2	E3	E4	E1	E2	E3	E4
P1	5	5	5	5	5	5	2	2	3	5	4	5	5
P2	4	3	3	5	2	—	—	—	—	4	2	5	3
P3	5	4	5	5	5	5	3	3	5	3	4	3	3
P4	5	4	5	4	5	5	5	5	4	5	3	2	3
P5	5	4	5	5	5	5	3	3	4	3	3	2	3
P6	5	1	3	4	5	3	1	2	3	4	3	3	3
P7	5	4	4	3	4	4	3	3	4	3	4	3	3
P8	5	5	5	5	5	4	3	3	4	3	4	3	3
P9	5	5	5	5	4	5	4	4	4	3	4	3	3
P10	5	4	5	5	4	5	5	4	4	5	5	5	4

Table 8.11: Experience items (1–5 Likert) per participant across all strategies. Strategy A: E1 Readable, E2 Comprehension, E3 Findability, E4 Information amount, E5 Trend visibility. Strategy B: E1 Readable, E2 Comprehension, E3 Findability, E4 Search quality. Strategy C: E1 Readable, E2 Overview, E3 Scannability, E4 Findability.

Modified SUS: Item Responses and Composite Scores

ID	Strategy A						Strategy B						Strategy C					
	S1	S2	S3	S4	S5	Σ	S1	S2	S3	S4	S5	Σ	S1	S2	S3	S4	S5	Σ
P1	5	1	5	1	5	100	–	–	–	–	–	–	5	1	5	1	2	85
P2	3	2	3	3	3	55	–	–	–	–	–	–	4	1	3	2	3	70
P3	5	1	5	1	4	95	4	2	4	2	4	75	3	2	3	2	3	60
P4	4	3	4	1	4	75	5	1	4	1	3	85	3	3	3	2	2	50
P5	5	1	5	1	3	90	5	3	3	2	2	60	3	4	3	4	2	35
P6	3	3	5	1	3	70	3	2	2	3	3	50	3	4	2	4	2	30
P7	2	2	2	4	4	45	3	3	4	3	3	55	4	4	4	4	3	50
P8	5	1	5	4	5	85	3	4	4	3	3	50	4	4	4	3	4	60
P9	4	3	4	2	5	75	4	2	5	2	5	85	4	1	3	4	4	65
P10	5	1	5	2	4	90	5	1	5	1	4	95	5	1	5	1	3	90

Table 8.12: Modified SUS item responses (S1–S5) and composite scores (Σ) per participant. S1: Overview (+), S2: Complexity (–), S3: Organization (+), S4: Clutter (–), S5: Confidence (+). Scoring: positive items contribute (response – 1), negative items contribute (5 – response), sum $\times$ 5.

Part D: Comparative Ratings

ID	Strategy A		Strategy B		Strategy C	
	Clarity	Ease	Clarity	Ease	Clarity	Ease
P1	4	4	3	3	4	4
P2	3	4	–	–	4	3
P3	5	5	4	5	3	3
P4	5	5	4	5	3	2
P5	5	5	4	3	3	3
P6	4	5	3	2	3	3
P7	5	5	4	3	3	3
P8	5	5	4	3	3	3
P9	5	5	3	4	2	2
P10	4	5	4	4	5	4

Table 8.13: Part D comparative ratings: data clarity and ease of use (1–5) per participant.

Part E: Ranking and Recommendation

ID	1st Choice	2nd Choice	3rd Choice	Rec. Score
P1	A	C	B	10
P2	C	A	B	6
P3	A	B	C	10
P4	A	B	C	8
P5	A	B	C	10
P6	A	C	B	8
P7	A	B	C	8
P8	A	B	C	8
P9	A	B	C	10
P10	C	A	B	10

Table 8.14: Overall ranking and recommendation score (1–10) per participant.

D Prototype Screenshots

Strategy A: Card Based Progressive Disclosure (Native iOS)




Figure 8.2: Strategy A detail views. Left: UAW detail sheet with tooltip and stacked bar chart. Center: Time scale filter with aggregation mode, time range presets, and date picker. Right: NFT category detail with ranked project table.

Strategy B: Narrative Contextualization (Figma)



Figure 8.3: Strategy B overview. Top row: Home screen with price chart, KPI row, and narrative text; Categories list; Core Metrics overlay with tooltip. Bottom row: Outline dApp detail; Core Metrics wallet breakdown; Outline UAW line chart with tooltip.




Figure 8.4: Strategy B detail views. Left: Gaming category transaction chart carousel. Center: Search overlay with recently searched items. Right: Search with keyboard input filtering results.



Figure 8.1: Strategy A overview. Top row: Home screen with stacked metric cards, Categories tab with donut chart, time scale filter sheet. Bottom row: Transactions detail (stacked area), Transaction Fees detail (stacked bar), Gas Fees detail (ranked horizontal bar).

Strategy C: Multi Chart Exploration (Figma)




Figure 8.5: Strategy C overview. Top row: Home screen with KPI grid; Categories icon grid; Categories ranked list with donut chart. Bottom row: Gaming Transactions (stacked area); Gaming UAW (stacked bar); Transaction Fees (radial bar chart).




Figure 8.6: Strategy C detail views. Left: Best performance treemap for Gaming wallet data. Center: Total UAW overview for Gaming. Right: Outmine Inflow/Outflow bidirectional bar chart.

E Mock Data Generation Parameters

The `MockAnalyticsRepository` generates all data locally using a linear congruential generator seeded with a fixed constant, producing deterministic output across launches. Asset price curves use geometric Brownian motion with Box-Muller transformed normal deviates, annual seasonality, and slow exponential moving average reversion. Metric time series combine per-dApp sinusoidal base patterns with randomized noise, then downsample to target point counts that vary by timeframe: 12 points for one day, 30 for one month, and 120 for all time.

Charts are rendered using Apple's Swift Charts framework. The line chart (`LineMark`) uses monotone interpolation and supports a drag gesture for continuous data point selection. The ranked chart for gas fees uses `GeometryReader` for proportional horizontal bar widths, with adaptive text placement that shifts labels outside the bar when the bar width falls below 35 percent of the available space.

F Contents of Digital Media

Directory	Files	Description
\Thesis	Bachelorarbeit_Garrels_Finn_1051281.pdf	The thesis in PDF format

Table 8.15: Contents of the digital media.

The source code for Strategy A and the Figma project files for Strategies B and C were developed in cooperation with FiDi Tech Inc. and are subject to confidentiality. They are not included on the digital medium but can be made available to the examiners upon request.

Eidesstattliche Erklärung

Hiermit erkläre ich an Eides statt, dass ich die vorliegende Arbeit mit dem Titel *Comparative Study of Blockchain Visualization Strategies on Mobile Devices* selbständig und ohne fremde Hilfe angefertigt habe. Es wurden keine anderen als die angegebenen Quellen und Hilfsmittel verwendet und alle wörtlich oder sinngemäß übernommenen Aussagen als solche kenntlich gemacht. Die Arbeit wurde weder vollständig noch in wesentlichen Teilen anderweitig für eine Prüfungsleistung verwendet oder eingereicht. Mir ist bekannt, dass ein Verstoß gegen diese Erklärung zum Nichtbestehen der Arbeit führen kann.

Osnabrück, _____

Unterschrift: _____