



UHASSELT

KNOWLEDGE IN ACTION



Maastricht University

Faculteit Wetenschappen **School voor Informatietechnologie**

master in de informatica

Masterthesis

A Multi-Model AI-Driven GUI Framework for Dynamic User Adaptation

Yarne Dirkx

Scriptie ingediend tot het behalen van de graad van master in de informatica

PROMOTOR :

Prof. dr. Kris LUYTEN

De transnationale Universiteit Limburg is een uniek samenwerkingsverband van twee universiteiten in twee landen: de Universiteit Hasselt en Maastricht University.



UHASSELT

KNOWLEDGE IN ACTION

www.uhasselt.be

Universiteit Hasselt

Campus Hasselt:

Martelarenlaan 42 | 3500 Hasselt

Campus Diepenbeek:

Agoralaan Gebouw D | 3590 Diepenbeek

2024
2025



Maastricht University

Faculteit Wetenschappen ***School voor Informatietechnologie***

master in de informatica

Masterthesis

A Multi-Model AI-Driven GUI Framework for Dynamic User Adaptation

Yarne Dirkx

Scriptie ingediend tot het behalen van de graad van master in de informatica

PROMOTOR :

Prof. dr. Kris LUYTEN

Abstract

This thesis presents a multimodal AI-driven GUI framework for dynamic user interface adaptation, enabling real-time, personalised accessibility enhancements across platforms such as Flutter, SwiftUI, and beyond. Designed for the health domain, the framework supports motor-impaired, visually impaired, and hands free users by integrating touch, keyboard, voice, and gestures. At its core is Smart Intent Fusion (SIF), a configurable multi-agent architecture powered by large language models (Google Gemini) that fuses inputs, events, profiles, and interaction history to infer intent and propose targeted adaptations (e.g., button enlargement, contrast enhancement, navigation mode switching). The reasoning combines rule-based logic for predictable, low-latency back-up responses with LLM-driven reasoning for complex or ambiguous cases. Developer integration is provided through standardised JSON contracts and a FastAPI backend. A feasibility study across six user profiles and 84 events shows that 97.5% of suggested changes are accessibility-oriented with high internal coherence ($DCI \approx 0.995$). In a balanced configuration, median latency is 13.2s with 84.5% schema-valid outputs and 55% profile-action alignment; a heavier configuration raises schema validity to 100% and alignment to 61%, at 36.1s median latency. Corrective responses are near-perfect for common motor/voice issues (missed taps, slider overshoots, voice commands) and weaker for gestures, indicating a clear path for improvement. The results demonstrate practical, accessibility-focused adaptation with measurable trade-offs and lay groundwork for future systems that can autonomously refactor UI code for richer, context-aware personalisation.

Acknowledgments

I would like to express my deepest appreciation to all those who have supported and guided me throughout the course of my master's thesis.

First and foremost, I am sincerely grateful to Prof. Dr. Kris Luyten, my thesis supervisor, for his invaluable guidance, constructive feedback, practical tips, and continuous support.

My heartfelt gratitude extends to my family, especially my sister Phaedra, my brother-in-law Dennis, and my parents for their unwavering support, patience, and help with proofreading.

I would also like to thank my fellow students for their encouragement, collaboration, and support during this journey. Special thanks go to my two dogs, Tobias and Casper, whose company during countless late-night coding sessions, and gentle reminders to take breaks helped keep me grounded.

Finally, I am thankful to my friends for their motivation, belief in my abilities, and for making this challenging journey a rewarding one.

This work would not have been possible without the contribution and support of all these people, to whom I am deeply grateful.

Nederlandse Samenvatting

Inleiding

Grafische gebruikersinterfaces (GUI's) vormen de kern van hoe we dagelijks met technologie omgaan, van smartphones tot websites en smartwatches. Deze interfaces maken complexe taken toegankelijk door visuele elementen zoals knoppen, iconen en menu's te gebruiken, in plaats van tekstgebaseerde commando's. Vóór de opkomst van GUI's in de jaren 1980 waren command-line interfaces (CLI's) dominant, wat computing beperkte tot experts die commando's moesten memoriseren. GUI's hebben dit veranderd door visuele metaforen te introduceren, waardoor technologie beschikbaar werd voor het brede publiek. Ze vertalen ingewikkelde computertaken naar intuïtieve, visuele acties.

In de loop der jaren zijn GUI's geëvolueerd van het klassieke WIMP-model (windows, icons, menus, pointers) naar multimodale systemen die touch, stem, gebaren en zelfs ruimtelijke interacties in augmented reality (AR) en virtual reality (VR) omvatten. Dit biedt rijkere en natuurlijkere interacties. Toch blijven de meeste GUI's statisch: ze volgen een one-size-fits-all aanpak, met dezelfde lay-out, gedrag en visuele elementen voor alle gebruikers, ongeacht hun behoeften of context. Dit creëert barrières in een diverse wereld, vooral voor gebruikers met toegankelijkheidsuitdagingen. Mensen met visuele beperkingen worstelen met laag contrast of gebrek aan screenreader-ondersteuning, terwijl motorisch beperkte gebruikers moeite hebben met kleine knoppen of precieze taps. Handsfree gebruikers, zoals in medische contexten, hebben alternatieve invoermethoden nodig.

Gebruikers interacteren bovendien in variërende omgevingen – lawaaiig, felverlicht of onderweg – wat interactie verder compliceert. Apparaten variëren in grootte en capaciteit, van smartwatches tot desktops, maar toegankelijkheidsfuncties zijn vaak inconsistent. Dit benadrukt een kernuitdaging: hoe ontwerpen we GUI's die dynamisch aanpassen aan individuele toegankelijkheidsbehoeften en context, voor inclusieve en efficiënte interactie?

Adaptieve gebruikersinterfaces (UI's) passen zich dynamisch aan aan gebruikersbehoeften en contextuele factoren. Ze wijzigen lay-out, gedrag of uiterlijk op basis van voorkeuren, vaardigheden en omstandigheden. Door multimodale invoer zoals touch, stem, gebaren en gaze te verwerken, bieden ze gepersonaliseerde, inclusieve ervaringen, vooral voor diverse gebruikersgroepen. Een 'intelligente' interface interpreteert complexe gedragingen autonoom en leert van interacties in real-time. Kunstmatige intelligentie (AI), inclusief machine learning en grote taalmodellen (LLM's), maakt dit mogelijk door invoer, UI-context en gebruikersdata te verwerken.

LLM's zoals GPT hebben natuurlijke taalverwerking gerevolutioneerd. Ze begrijpen context, leiden intentie af en genereren menselijke reacties, ideaal voor multimodale interpretatie. In adaptieve UI's fungeren LLM's als controllers die interacties vertalen naar aanpassingen, zoals het vergroten van knoppen op basis van een stemcommando of aarzeling in gebaren. Ze kunnen zelfs UI-elementen genereren door de interface als API te behandelen.

Ondanks onderzoek en technologie blijven adaptieve UI's zeldzaam door implementatiecomplexiteit en gebrek aan frameworks. Dit motiveerde deze scriptie, geïnspireerd op ideeën om UI's als API's te behandelen, bestuurd door intelligente modellen. LLM's bieden potentieel, maar hebben uitdagingen zoals latency, hallucinaties en gebrek aan specialisatie voor GUI-taken. Deze scriptie ziet LLM's als contextbewuste assistenten die, met begeleiding, UI-aanpassingen verbeteren.

Probleemstelling en Onderzoeksdoelen

De meeste applicaties zijn statisch en passen zich niet aan veranderende contexten of beperkingen. Huidige adaptaties zijn cosmetisch, zoals aanpassing aan schermgrootte, maar pakken niet de diepere uitdagingen aan van intentie, cognitieve staat of fysieke limieten in real-time. Het ontwerpen van adaptieve UI's is ongewoon door technische moeilijkheden in modellering en gebrek aan frameworks. Real-world interactie is chaotisch: kleine variaties zoals een miss-tap of omgevingsruis leiden tot divergenties, wat regelgebaseerd ontwerp fragiel maakt.

Dit vereist systemen die nuances interpreteren, geschikt voor AI zoals LLM's. Bestaande frameworks integreren zulke modellen zelden voor real-time adaptatie, wat een kloof creëert tussen potentieel en praktijk. Het primaire doel is een modulair, multimodaal AI-framework te ontwerpen, implementeren en valideren voor dynamische UI-aanpassing, met focus op toegankelijkheid. Het gaat voorbij statische designs met real-time, contextbewuste aanpassingen via hybride redenering.

Specifieke doelen:

1. Modulaire cross-platform architectuur ontwikkelen voor events over Flutter, SwiftUI en meer.
2. Multimodale invoerfusie implementeren (touch, stem, gebaren).
3. Smart Intent Fusion (SIF) integreren: hybride engine voor aanpassingen.
4. Toegankelijkheidsgerichte aanpassingen leveren (vergroten, contrast, modi).
5. Ontwikkelaarvriendelijk integratiepad bieden via een soort basic SDK.
6. Performance evalueren via proof-of-concept en metrics.

Gerelateerd Werk

Multimodaal AI in gebruikersinterfaces (UI's) omvat een evolutie van pointing devices, van traditionele muisinteracties naar touch-gebaseerde systemen op mobiele apparaten, die precisie en intuïtie verbeteren maar uitdagingen stellen voor gebruikers met beperkingen. Steminterfaces, zoals Siri en Alexa, introduceren natuurlijke taalverwerking (NLP) voor handsfree bediening, terwijl gebaren en gaze-tracking – geïllustreerd door Kinect voor lichaamstracking en eye-tracking tools voor cursorcontrole – rijkere, niet-contact interacties mogelijk maken, vooral in AR/VR-contexten.

Adaptieve GUI's richten zich primair op toegankelijkheid, met systemen zoals SUPPLE (dat constraint-optimalisatie gebruikt voor motorische calibratie en lay-outgeneratie) en GUIDe (dat gaze combineert met keyboard of muis voor nauwkeurige targeting en scrolling). Uitdagingen liggen in het handhaven van consistentie over modaliteiten, waar variërende inputs (bijv. touch vs. stem) kunnen leiden tot inconsistente ervaringen, en in het balanceren van automatisering zonder gebruikerscontrole te verliezen.

Klassieke adaptieve technieken includeren responsive layouts, die media queries en breakpoints gebruiken voor apparaatadaptatie (bijv. aanpassing aan schermgrootte en oriëntatie), en context-aware design dat sensoren (locatie, licht, beweging) integreert voor omgevingsaanpassingen, zoals thema-switches in donkere modi. Programmeerbare UI's en frameworks bieden meer dynamiek: Reflow analyseert pixels voor runtime lay-outwijzigingen, UICoder genereert code met LLM's voor nieuwe UI-creatie, en RL-gebaseerde adaptatie (reinforcement learning) leert van interacties om lay-outs en content te optimaliseren, vaak met gespecialiseerde hardware.

Gebruikersprofielgebaseerde UI's, zoals XML-gebaseerde runtime systemen, maken dynamische beschrijvingen mogelijk, aangepast aan mobiele en embedded devices voor personalisatie op basis van voorkeuren en capaciteiten. Multimodale fusie omvat architecturen zoals vroeg (feature-level) of laat (semantisch-level) fusie, met eventmodellen die inputs standaardiseren voor coherente verwerking.

LLM's als UI-controllers transformeren interfaces door ze als API's te behandelen, waar modellen intentie afleiden en acties uitvoeren; agents zoals ReAct (redeneren en handelen in loops) en AutoGPT (autonome taakuitvoering) voegen autonomie toe. In gezondheidstoepassingen tonen chatbots voor cognitieve gedragstherapie (bijv. Woebot voor depressie) of apps voor rugpijnbehandeling potentieel voor gepersonaliseerde, multimodale zorg, met integratie van stem en gebaren voor therapie-ondersteuning. Dit werk bouwt hierop door multimodale fusie te combineren met hybride AI (regels + LLM's) voor

real-time, toegankelijkheidsgerichte adaptatie, overbruggend gaps in bestaande systemen zoals statische generatie of beperkte modaliteiten.

Systeemontwerp en Architectuur

Het framework is opgebouwd rond een drie-laagse architectuur die modulariteit en schaalbaarheid prioriteert, met een duidelijke scheiding van verantwoordelijkheden om uitbreidbaarheid te garanderen. De lagen zijn: de Frontend laag voor directe UI-interactie, de Input Adapter laag voor multimodale invoerverwerking, en de SIF Backend laag voor intentiefusie en adaptatiebeslissingen. Het ontwerp focust primair op het gezondheidsdomein, waar het ondersteunt gebruikers met motorische beperkingen (bijv. tremors of coördinatieproblemen), visuele beperkingen (bijv. laag zicht of kleurenblindheid) en hands-free behoeften (bijv. in revalidatie of mobiele contexten), door aanpassingen te bieden die inclusiviteit bevorderen zonder de kernfunctionaliteit te verstoren.

De Frontend laag, geïmplementeerd in frameworks zoals Flutter of SwiftUI, beheert de interface-elementen zoals knoppen, sliders, toggles en navigatiecomponenten. Het ondersteunt aanpassingsniveaus op visueel vlak (bijv. kleurcontrast verhogen of fonts vergroten) en interactief vlak (bijv. invoermodi schakelen van touch naar stem). Deze laag vangt ruwe interactiegebeurtenissen op en past aanpassingen toe met soepele animaties voor een naadloze gebruikerservaring, terwijl het ook profielinitialisatie handhabet via een editor voor persoonlijke instellingen.

De Input Adapter laag fungeert als intermediair en standaardiseert diverse invoermodaliteiten – touch (taps en swipes), keyboard (toetsaanslagen), stem (spraakherkenning via API's) en gebaren (via sensoren zoals accelerometer) – in een uniform JSON-formaat met velden zoals type, timestamp en confidence-score. Dit zorgt voor consistente verwerking, onafhankelijk van het platform. Transport wordt beheerd via WebSocket voor real-time, lage-latency events (ideaal voor continue interactie) en HTTP voor batch-operaties zoals profielupdates, met automatische reconnects voor robuustheid.

De SIF Backend laag integreert alles door data te fuseren: het combineert gestandaardiseerde events met profielcontext using regelgebaseerde logica voor snelle, deterministische reacties (bijv. automatische knopvergroting bij detectie van een miss-tap) en LLM-redenering voor complexere gevallen (bijv. ambiguë stemcommando's interpreteren). Gebruikersprofielen modelleren beperkingen (impairments), voorkeuren (bijv. hoge contrastmodus) en interactiegeschiedenis (logs van eerdere events) voor diepgaande personalisatie, opgeslagen in een database zoals MongoDB voor persistentie.

Dynamische aanpassingsmechanismen volgen een pipeline: inkomende events worden verwerkt, intentie afgeleid (bijv. via heatmap-analyse van tap-patronen), en aanpassingen gegenereerd (bijv. knopvergroting met factor 1.5 of contrastverhoging naar 4:1 ratio, in lijn met WCAG-richtlijnen). Een continue leerlus integreert feedback door geschiedenis te analyseren en modellen te verfijnen, zodat aanpassingen evolueren (bijv. progressieve optimalisatie bij herhaalde fouten). Ontwerpoverwegingen benadrukken modulariteit (losse koppeling voor eenvoudige uitbreidingen, bijv. nieuwe modaliteiten toevoegen), schaalbaarheid (asynchrone verwerking voor hoge loads) en privacy (data-redactie op device-niveau, minimale logging van gevoelige inputs), met toekomstige extensies naar VR/AR via Unity. Dit zorgt voor een robuust, developer-vriendelijk systeem dat statische UI's omzet in adaptieve, contextbewuste interfaces.

Smart Intent Fusion (SIF)

Smart Intent Fusion (SIF) vormt het hart van het framework en is verantwoordelijk voor het intelligent combineren van multimodale invoer met gebruikerscontext om intenties af te leiden en gepersonaliseerde UI-aanpassingen voor te stellen. In essentie is SIF een geavanceerde redeneerengine die multimodale signalen – zoals touch-events, stemcommando's, gebaren en toetsaanslagen – fuseert met profielgegevens, interactiegeschiedenis en UI-context. Dit resulteert in real-time aanpassingen die de toegankelijkheid verbeteren, zoals het vergroten van knoppen voor motorisch beperkte gebruikers of het verhogen van contrast voor visueel beperkten. SIF is ontworpen als een multi-agent architectuur, aangedreven door grote taalmodellen (LLM's) via de Google Gemini API, maar met een hybride aanpak om de beperkingen van LLM's te compenseren.

Theoretisch is SIF geworteld in multimodale fusieprincipes uit de HCI-literatuur. Multimodale fusie kan vroeg (op feature-niveau, bijv. ruwe data van sensoren combineren) of laat (op semantisch niveau,

bijv. geïnterpreteerde intenties integreren) gebeuren. SIF kiest voor een late fusie om complexiteit te beheren: invoer wordt eerst gestandaardiseerd in een uniforme JSON-structuur voordat ze worden verwerkt. Intentie-inferentie bouwt op modellen zoals ReAct (Reasoning and Acting), waarbij LLM's redeneren over context om acties te voorspellen. De hybride aanpak – regelgebaseerde logica voor eenvoudige, voorspelbare gevallen en LLM-redenering voor ambigue of complexe scenario's – zorgt voor betrouwbaarheid. Regels bieden lage latency (milliseconden), terwijl LLM's diepere inzichten leveren, maar met mogelijke vertraging (seconden). Dit hybride model voorkomt over-reliance op LLM's, die vatbaar zijn voor hallucinaties (fictieve outputs) of inconsistentie.

Gebruikersprofielen spelen een cruciale rol in SIF en structureren behoeften op een gestandaardiseerde manier. Een profiel bevat velden zoals "impairments" (bijv. "motor" voor tremors, "visual" voor laag zicht), "preferences" (bijv. voorkeur voor grote fonts of stemfeedback) en "history" (een log van eerdere interacties). Deze profielen beïnvloeden beslissingen door als context te dienen: voor een motorisch beperkt profiel prioriteert SIF aanpassingen zoals knopvergroting of tremor-correctie. Leren van interactiegeschiedenis gebeurt continu; SIF analyseert patronen, zoals herhaalde miss-taps, om toekomstige aanpassingen te verfijnen. Dit creëert een feedbacklus waarbij het systeem personaliseert op basis van evoluerende data, in lijn met ability-based design principes uit HCI, waar interfaces zich aanpassen aan individuele capaciteiten in plaats van een standaard te forceren.

De modellering van multimodale invoerfusie begint met standaardisatie: ruwe events (bijv. een touch-tap met coördinaten en timestamp) worden omgezet in een JSON-formaat met velden zoals "type" (touch/voice/gesture), "confidence" (een score van 0-1 voor betrouwbaarheid, bijv. spraakherkenning accuracy) en "timestamp". Timing is essentieel; SIF gebruikt vensters (bijv. 500ms) om synchrone inputs te groeperen, zoals een stemcommando tijdens een gebaar. Confidence-scores wegen inputs: een lage confidence-touch (bijv. door tremor) triggert extra LLM-redenering. LLM's worden ingezet voor fusiebeslissingen, bijvoorbeeld door prompts zoals "Gegeven deze events en profiel, leid intentie af en stel aanpassingen voor", wat resulteert in outputs zoals {"action": "enlarge.button", "target": "submit", "value": 1.5}.

Regelgebaseerde logica handelt eenvoudige gevallen: bijv. als een miss-tap wordt gedetecteerd (afstand > threshold), vergroot de knop automatisch. Dit is deterministisch en laag-latency, ideaal voor real-time. Voor ambiguïteit schakelt SIF naar LLM's, die context interpreteren (bijv. "gebruiker aarzelt bij slider – stel voice-control voor"). Heatmap-analyse voegt toe door interactiepatronen te visualiseren: herhaalde taps rond een knop genereren een heatmap, die SIF gebruikt om hotspots te identificeren en lay-outs aan te passen.

Multi-Agent SIF (MA-SIF) verfijnt dit door taken te splitsen in gespecialiseerde agents: de UI Agent focust op visuele aanpassingen (bijv. contrast, lay-out), de Geometry Agent op ruimtelijke veranderingen (bijv. knopgrootte, positie), en de Input Agent op modaliteitsswitches (bijv. van touch naar voice). Elke agent ontvangt dezelfde input maar met een specifieke prompt, en genereert onafhankelijke voorstellen. Een Validator Agent reconcileert outputs: het controleert op conflicten (bijv. overlappende acties), valideert tegen een strikt JSON-schema (bijv. "value" moet numeriek zijn binnen ranges), en mergeert tot een finale set. MA-SIF is configureerbaar via een JSON-file, waar agents' prompts, toegestane acties en temperaturen (creativiteitsparameter, laag voor conservatief) worden gedefinieerd. Bijvoorbeeld: temperature=0.2 voor deterministische outputs, thinking_budget=500 tokens om redenering te beperken.

Prompt engineering is cruciaal voor LLM-prestaties. Principes includeren expliciete instructies ("Wees conservatief, hallucineer geen targets"), gestructureerde formats (bijv. "Output als JSON: actions: [type, target, value]"), en ambiguïteit vermijden (geen disjuncties zoals "A of B" die LLM's verkeerd interpreteren). Strikte schemas dwingen validiteit af: ongeldige outputs worden verworpen en fallbacken naar regels. Metrics evalueren SIF: latency (mediaan 13s voor MA-SIF), schema-validiteit (84.5% in tests), correctheid (actie-alignment met profiel, 97.5% toegankelijkheidsgericht), en personalisatie (bijv. % aanpassingen uniek per profiel).

Beperkingen van LLM-integratie zijn significant: latency door API-calls (seconden vs. ms voor regels), hallucinaties (bijv. niet-bestaande targets), en tokenlimits (contextgrootte beperkt geschiedenis). Mogelijke oplossingen: hybride fallbacks (regels eerst, LLM asynchroon), validatoren voor filtering, en caching van bijvoorbeeld veelvoorkomende prompts. LLM-selectie (Gemini vs. GPT) balanceert kosten, snelheid en multimodaliteit. Toekomstige richtingen: autonome on-device agents met fine-tuned modellen (bijv. getraind op UI-logs en WCAG-richtlijnen), integratie met RL voor leren van feedback, en visuele UI-analyse (vision models voor screenshot-parsing) om context te verrijken zonder metadata. SIF

transformeert statische UI's naar dynamische, intentiebewuste systemen, met een balans tussen snelheid, intelligentie en betrouwbaarheid, cruciaal voor toegankelijkheidsgerichte toepassingen.

Implementatie van het Framework

De implementatie van het framework is gerealiseerd als een proof-of-concept: een Adaptieve Smart Home Controller, een eenvoudige app die slimme apparaten (lamp, thermostaat, slot) beheert en dynamisch aanpast op basis van gebruikersevents. Dit demonstreert de end-to-end workflow en portability over platforms. De stack omvat Flutter voor de frontend (Dart-based, cross-platform UI), een Dart-adapter voor inputverwerking, en een Python-backend met FastAPI voor SIF-logica. Ontwikkelomgeving: Flutter SDK voor UI, Dart voor adapter, Python 3 met libraries zoals FastAPI, MongoDB-driver en Google Gemini API-client.

De Frontend in Flutter beheert de UI-staat en interacties. Het bestaat uit scrollbare kaarten met minimalistische controls (toggle, slider, knop) en een mock-event rij voor testen (bijv. "Miss Tap", "Voice Command"). Events worden vastgelegd (bijv. tap-coördinaten, stemtekst) en doorgestuurd naar de adapter. Aanpassingen worden toegepast via een state-model: bij ontvangst van JSON-outputs (bijv. {"action": "enlarge_button", "target": "submit", "value": 1.5}) update de UI dynamisch met animaties (bijv. AnimatedContainer voor soepele resizing). State wordt beheerd met Provider voor reactiviteit, en profielen bootstrappen via een editor-scherm waar gebruikers impairments selecteren. Voor responsiveness: een loading-indicator (roterende gradient) tijdens backend-wachting, en partial results (regels eerst, LLM later).

De Input Adapter in Dart fungeert als brug: het serialiseert ruwe events naar JSON, beheert transport (WebSocket voor real-time events, HTTP voor profielbeheer), en callbackt aanpassingen terug naar de frontend. Klasse-overzicht: AdaptiveUIAdapter met methods zoals sendEvent() en onAdaptationReceived(). Interne representaties: Event (type, data, timestamp), Adaptation (action, target, value), UserProfile (impairments, preferences). Extensibiliteit: nieuwe modaliteiten (bijv. gaze) toevoegen door Event-subklassen. WebSocket-lifecycle includes reconnect-backoff voor robuustheid.

De SIF Backend in Python gebruikt FastAPI voor endpoints (/events via WebSocket, /profiles via HTTP) en MongoDB voor persistentie (geschiedenis opslaan per user_id). SIF/MA-SIF wordt geïmplementeerd als async functions: events verrijken met profiel/geschiedenis, regels toepassen (bijv. if miss-tap: enlarge), dan LLM-invoatie (Gemini met prompts uit config). Structured outputs via JSON-schemas, guardrails tegen invaliditeit (validator parseert en filtert). Regel-fallbacks: hardcoded triggers (bijv. overshoot-slider → voice-mode). Heatmap-analyse: accumuleer taps in een grid, detecteer clusters voor lay-out shifts. Latency-handling: timeouts (5s), partial results (stuur regels direct). Security: CORS voor frontend-origin, geen auth in prototype.

Profielen worden bootstrapped (default laden) en bewerkt via HTTP, met opslag in MongoDB voor persistentie. Dynamische mechanismen: applyAdaptations() mapped acties naar widgets (bijv. enlarge → scale factor), met animaties voor UX. Conflicten: prioriteer (bijv. geometry > UI), onbekende acties negeren. Real-time voorbeeld: miss-tap → regel-vergroting + LLM-suggestie (contrast boost), toegepast na 13s mediaan. Backend-injectie interface (web-app) simuleert events/profielen, visualiseert responses en geschiedenis voor debugging. Cross-platform voorbeeld in SwiftUI: minimal UI met state-mapping, adapter via URLSession/WebSocket, events injecteren, toont portability met dezelfde JSON-contracts.

Ontwerpbeslissingen benadrukken modulariteit (losse koppeling voor swaps), WebSocket voor lage-latency events (vs. HTTP batch), MongoDB voor flexibele schema's (JSON-docs), en hybride redenering (regels voor basis, MA-SIF voor diepte). Uitdagingen: LLM-consistentie opgelost met validatoren en prompts; performance met caching en batching; beveiliging met CORS en toekomstige auth (JWT). Testen met incomplete modaliteiten via mocks; privacy via on-device redactie.

Deze implementatie bewijst de haalbaarheid: een schaalbaar, developer-vriendelijk framework dat statische UI's omzet in adaptieve, toegankelijke systemen, met focus op gezondheidstoepassingen.

Haalbaarheidsstudie

De evaluatie volgt een haalbaarheidsstudie om te verifiëren of de multimodale AI-gedreven adaptatiepijn end-to-end werkt onder realistische interactiesporen en of dit nuttig is voor toegankelijkheid. Zes uiteenlopende profielen werden gebruikt (motorisch, visueel, handsfree), elk met twee opeenvolgende runs

van zeven events (totaal 84): `miss_tap`, `slider_miss`, `voice` en `gesture`. Dit bootst real-world gebruik na door directe fouten te combineren met een groeiende geschiedenis, zodat de tweede run voortbouwt op de eerste voor contextueel leren.

De MA-SIF-configuratie (multi-agent, gebalanceerde instellingen) is getoetst op kernmetrics: schema-validiteit (84,52% van de responsen voldoet aan het strikte JSON-schema; 100% passeert via de Validator), toegankelijkheidsaandeel (97,51% van alle acties is expliciet op toegankelijkheid gericht), en interne samenhang ($DCI \approx 0,995$). Latency in deze configuratie heeft een mediaan van 13,19 s; regels leveren directe feedback, terwijl LLM-adaptaties asynchroon binnenkomen. Qua geschiktheid per eventtype is ERA zeer hoog voor motor-gerelateerde fouten (`miss_tap` en `slider_miss` beide 100,00%), hoog voor `voice` (97,22%) en laag voor `gesture` (8,33%). Per profiel varieert schema-validiteit van 71,43% (P3, handsfree) tot 100,00% (P5, visueel+motor).

De discussie benadrukt sterke personalisatie – aanpassingen evolueren met geschiedenis (bijv. herhaalde miss-taps leiden tot progressieve vergrotingen) – en tegelijk de beperkingen van de synthetische opzet: geen echte gebruikers betekent beperkte gedragsvariatie en geen subjectieve feedback; latency is merkbaar in snelle interacties. Conclusie: het systeem is haalbaar voor toegankelijkheidsverbetering met hoge validiteit en alignment, maar vraagt vervolgwerk via gebruikersstudies en gerichte optimalisatie van prompts, validatorbeleid en fallbacks.

Discussie en Toekomstig Werk

De implicaties reiken tot inclusieve HCI: statische interfaces verschuiven naar dynamische systemen die zich tijdens gebruik aanpassen aan uiteenlopende behoeften, met duidelijke relevantie voor gezondheidsccontexten (bijv. revalidatie-apps of stemgestuurde tools). Belangrijkste bijdragen zijn een modulaire drie-laagse architectuur (Frontend, Input Adapter, SIF Backend), een hybride redeneringskern (regels + multi-agent LLM's met Validator) en een focus op concrete toegankelijkheidsaanpassingen (knopvergroting, contrast, moduswissels) via ontwikkelaarsvriendelijke JSON-contracten.

In vergelijking met gerelateerd werk onderscheidt dit framework zich door multimodale fusie en runtime-adaptatie: t.o.v. SUPPLE (compile-time layout-optimalisatie met calibratie) biedt het runtime-fusie van signalen; t.o.v. Reflow (pixel-gebaseerde layout) voegt het semantische intentieredenering toe; UICoder (LLM-codegeneratie) werkt compile-time, terwijl dit werk bestaande live UI's aanpast; RL-gebaseerde systemen vragen vaak gespecialiseerde hardware en langere training.

Toekomstig werk omvat een gebruikersstudie met 12–18 deelnemers (4–6 per cluster) in een within-subjects design (regels-only vs. regels+MA-SIF), met SUS, NASA-TLX, foutenpercentages en taaktijden in smart home-scenario's met geïnduceerde beperkingen. Verdere stappen zijn overlays/OS-hooks voor retrofit zonder codewijzigingen, een ontwikkelaars-SDK (Flutter/SwiftUI/Unity) met debug-tools en WCAG-templates, en een UI-analyzer die widget trees of screenshots semantisch begrijpt (vision-modellen) om context te verrijken. Uitbreiding van input omvat gaze-tracking, geavanceerde handtracking en BCI. Autonome agents kunnen prompts en beleid dynamisch bijsturen en zichzelf valideren; een gespecialiseerd, op logs en WCAG getraind on-device model verlaagt latency en verbetert privacy. Langetermijn: integratie van RL voor continue optimalisatie met dual-mode adaptatie (compile-time defaults + runtime tweaks).

Conclusie

Dit werk ontwikkelt een framework voor (nagenoeg) real-time multimodale UI-adaptatie met een sterke toegankelijkheidsfocus. Smart Intent Fusion (SIF) combineert deterministische regels met LLM-gedreven redenering in een multi-agent opzet met Validator en levert betrouwbare, contextbewuste aanpassingen (bijv. elementen vergroten of naar stemmodus schakelen) binnen een modulaire, platformoverschrijdende architectuur.

Bijdragen omvatten gestandaardiseerde JSON-contracten, ontwikkelaarsvriendelijke integratie en een hybride engine die snelheid (regels, snelle paden) en intelligentie (LLM's + Validator) in balans brengt. Beperkingen zijn afhankelijkheid van externe LLM-API's, een synthetische evaluatie (geen veldstudie) en merkbare latency in multi-agent flows (mediaan 13,19 s). De evaluatie bevestigt haalbaarheid: 84,52% schema-validiteit, 97,51% toegankelijkheidsgerichte acties en hoge interne samenhang ($DCI \approx 0,995$) over

84 events en zes profielen; motor- en voice-cases scoren zeer hoog, terwijl gebaren de grootste verbetermogelijkheid vormen. Een regels-only aanpak presteert niet beter, maar een rule fast path is praktisch (ca. 86% van de events heeft minstens één canonicke correctie): emit die onmiddellijk en laat MA-SIF + Validator de volledige, profielbewuste set consolideren. Daarmee vormt dit werk een solide basis voor verdere validatie met gebruikers en voor uitbreiding richting UI-bewuste context en on-device redenering.

Contents

1	Introduction	21
1.1	Background and Motivation	21
1.2	Problem Statement	22
1.3	Research Objectives	23
1.4	Thesis Structure	24
2	Related Work	25
2.1	Multimodal AI in User Interfaces	25
2.1.1	Pointing Devices and Touch Interfaces	25
2.1.2	Voice-Driven Interfaces	26
2.1.3	Gesture and Gaze Integration	26
2.2	Adaptive GUIs Across Modalities and Platforms	26
2.2.1	Accessibility-Focused Adaptive GUIs (Desktop & Mobile)	27
2.2.2	Design Challenges for Adaptive GUIs Across Modalities	27
2.3	Classical Adaptive UI Techniques	27
2.3.1	Responsive Layouts	28
2.3.2	Context-Aware Design	28
2.4	Programmable UIs	28
2.4.1	Reflow	28
2.4.2	UICoder	29
2.4.3	User Interface Adaptation using Reinforcement Learning	29
2.5	User Profile built UIs	30
2.5.1	XML-Based Runtime UI Systems	30
2.6	Multimodal Fusion and Input Event Modeling	31
2.6.1	Fusion Architectures	31
2.6.2	Event Abstraction Models	31
2.7	LLMs as UI Controllers	32
2.7.1	Turning UIs into APIs	32
2.7.2	Agents	32
2.8	Health and Accessibility Applications	33
3	System Design and Architecture	35
3.1	Introduction to System Design	35
3.2	Overview of the System Architecture	35
3.2.1	Accessibility Focus and Target User Groups	37
3.2.2	Accessibility Grounding and Adaptation Rationale (WCAG 2.1/2.2)	38
3.3	Frontend Layer: UI Design and Interaction	38
3.3.1	Interface Elements	38
3.3.2	Adaptation Levels	39
3.4	Input Adapter Layer: Multimodal Input Processing	39
3.5	SIF Backend Layer: Smart Intent Fusion (SIF)	40
3.6	User Profiles and Context Modeling	41
3.7	Dynamic Adaptation Mechanisms	41
3.7.1	Adaptation Pipeline	41
3.7.2	Supported Adaptation Actions	42
3.7.3	Continuous Learning and Feedback Loop	42

3.7.4	Design Considerations	42
3.8	Chapter Summary	42
4	Smart Intent Fusion (SIF)	45
4.1	Introduction to Smart Intent Fusion	45
4.2	Theoretical Foundations of Smart Intent Fusion	46
4.2.1	Multimodal Fusion	47
4.2.2	Intent Inference	47
4.2.3	Why Hybrid Works Best	48
4.2.4	Connection to Accessibility	48
4.3	User Profile and Context Integration	48
4.3.1	User Profiles	48
4.3.2	User Profile Structure	49
4.3.3	How Profiles Affect Decisions	49
4.3.4	Continuous Learning from History	50
4.3.5	Role in Accessibility	50
4.4	Modeling Multimodal Input Fusion	51
4.4.1	Event Standardisation	52
4.4.2	Timing and Confidence	52
4.4.3	LLM Reasoning in Fusion Decisions	53
4.5	Rule-Based Logic and LLM-Driven Adaptation	53
4.5.1	Rule-Based Logic	53
4.5.2	LLM-Driven Reasoning	54
4.5.3	Hybrid Approach in SIF	54
4.5.4	Heatmap Analysis	55
4.6	Multi-Agent Smart Intent Fusion (MA-SIF)	55
4.6.1	Why Multiple Agents?	55
4.6.2	Agent Roles	55
4.6.3	Adaptation Flow	56
4.6.4	Dynamic Configuration	57
4.6.5	Temperature and Thinking Budget	58
4.6.6	Example in Action	58
4.6.7	Benefits of the Multi-Agent Approach	59
4.7	Prompt Engineering for LLMs in SIF	59
4.7.1	LLM Prompt Design Principles	59
4.7.2	Prompt Structure from <code>sif_config.json</code>	60
4.7.3	Disjunction Ambiguity in LLM Interpretation	60
4.7.4	Balancing Model Parameters	61
4.7.5	Avoiding Hallucinations and Bad Values	61
4.7.6	Importance of a Strict JSON Schema	61
4.8	Performance and Evaluation Metrics for AI Logic	62
4.9	Limitations and Solutions of LLM Integration	65
4.9.1	LLM selection	65
4.9.2	Reliability and Latency Constraints	65
4.9.3	Hallucinations and Invalid Output	65
4.9.4	Token Limits and Context Size	66
4.9.5	Validator Complexity	66
4.9.6	Dependency on External APIs	66
4.10	Future Directions for AI-Driven Adaptation	66
4.11	Chapter Summary	67
5	An Adaptive Multimodal GUI Framework using LLMs	69
5.1	Introduction to an Adaptive Smart Home Controller	69
5.2	Development Environment	70
5.3	Frontend (Flutter): Adaptive Smart Home Controller	70
5.3.1	Responsibilities & Data Flow	71
5.3.2	UI Composition & State Model	71
5.3.3	Event & Adaptation Contract	72
5.3.4	Responsiveness & Feedback	73

5.3.5	Profile Bootstrap & Editing	73
5.3.6	Testing Harness: Mock Events & Input Capture	74
5.3.7	Summary	74
5.4	Input Adapter (Dart): Transport, Serialization & Adaptation Callback	75
5.4.1	Class Overview	75
5.4.2	Internal Representations of Event, Adaptation and User Profiles	75
5.4.3	Transport & Profile Management	76
5.4.4	Extensibility Example	77
5.4.5	Summary	77
5.5	SIF Backend Layer: Implementation of Adaptation Logic	78
5.5.1	Webserver layout and endpoints	78
5.5.2	Data persistence and history management	78
5.5.3	Smart Intent Fusion and MA-SIF	78
5.5.4	LLM invocation (Gemini)	79
5.5.5	Structured outputs and guardrails	80
5.5.6	Rule-based fallback and resilience	80
5.5.7	Heatmap Analysis	80
5.5.8	Latency, partial results, and error handling	81
5.5.9	Security and CORS considerations	81
5.5.10	Summary	81
5.6	User Profile and Context Implementation	81
5.7	Dynamic Adaptation Mechanisms Implementation	82
5.7.1	Application Mechanics (State, Animation, Ordering)	82
5.7.2	Conflicts and Unknown Actions	83
5.7.3	Real-Time Adaptation Example	83
5.8	Backend Injection Interface	85
5.8.1	Purpose and Scope	85
5.8.2	Architecture and Data Flow	85
5.8.3	Controls: Profiles and Events	85
5.8.4	Adaptation Response View	85
5.8.5	Interaction History Panel	85
5.8.6	Operational Notes and Limitations	86
5.9	Cross-Platform SwiftUI Example	86
5.9.1	Purpose and Scope	86
5.9.2	Adapter and Transport	86
5.9.3	Minimal UI and State Mapping	86
5.9.4	Event Injection	86
5.9.5	Usage and Limitations	87
5.10	Design Decisions	87
5.10.1	Modularity Over Monolithic Design	87
5.10.2	WebSocket for Real-Time vs. HTTP for Batch Processing	87
5.10.3	MongoDB for Persistent Storage	88
5.10.4	Rule-Based Fallback with LLM Reasoning	88
5.10.5	Multi-agent LLM reasoning (MA-SIF) vs single-agent LLM reasoning (normal SIF)	88
5.11	Implementation Challenges and Solutions	89
5.11.1	LLM reliability and output consistency	89
5.11.2	Performance under real-time constraints	89
5.11.3	Safeguards against malicious or replay attacks	89
5.11.4	Testing with incomplete modalities	89
5.11.5	Security and trust boundaries	89
5.12	Chapter Summary	90
6	Feasibility Study	91
6.1	Study Overview	91
6.2	Methodology	91
6.3	Accessibility and Design Quality	93
6.3.1	Profile–Action Alignment (PAA)	93
6.3.2	Error to Response Appropriateness (ERA)	93

6.3.3	Mode Enablement for Hands free (MEH)	93
6.3.4	Design Coherence Index (DCI)	94
6.3.5	Global accessibility share and WCAG coverage	94
6.3.6	Motor Benefit Proxy (Fitts law, analytic)	94
6.3.7	Appropriateness by event type	95
6.3.8	Stability across runs	95
6.3.9	Specificity via profile swap	96
6.3.10	Failure taxonomy (invalid cases)	96
6.3.11	Rule-only sufficiency (proxy)	97
6.3.12	Rule baseline vs. LLM (minimal vs. maximal)	97
6.3.13	Micro-cases on P1 and P5 profiles	98
6.4	Adaptation Performance (Latency)	99
6.4.1	Overall summary	99
6.4.2	Per profile outcomes	99
6.4.3	Configuration comparison (all profiles)	100
6.5	Discussion of Results	100
6.6	Study Limitations	101
6.7	Conclusion	101
7	Discussion and Future Work	103
7.1	Overview	103
7.2	Implications for Accessibility and HCI	103
7.3	Key Findings and Contributions	103
7.4	Comparison with Related Work	104
7.5	Future Work	105
7.5.1	Short-Term Improvements	105
7.5.2	Research Directions	106
8	Conclusion	107
8.1	Summary of the Work	107
8.2	Summary of Contributions	108
8.3	Limitations	108
8.4	Lessons Learned	108
8.5	Evaluation Results: Conclusion	109
8.6	Self-Reflection and Reflection on the Work	109
8.7	Final Remarks	110

List of Figures

3.1	High-level architecture flow: frontend events are standardised and sent to the backend, where reasoning fuses user context and generates adaptation outputs for the UI.	36
3.2	High-level flow of the Input Adapter Layer	40
3.3	Adaptation pipeline flow.	42
4.1	Abstract event flow in Smart Intent Fusion: events from the Input Adapter are enriched with profile/context, processed by rule-based and LLM/MA-SIF reasoning, and returned as adaptation outputs to the frontend.	46
4.2	Multi-agent SIF adaptation flow: user events and profile/history are processed by specialised agents, validated, and merged into final adaptations. Fallbacks ensure robustness.	56
5.1	Frontend data flow: events to adapter; adaptations back to UI.	71
5.2	Smart Home Controller UI: scrollable device cards (Lamp, Thermostat, Lock) with minimalist controls and the mock event row (<i>Miss Tap, Voice Command, Gesture</i>).	71
5.3	Loading indicator: rotating gradient card border while waiting for backend.	73
5.4	Profile editor for accessibility needs and UI preferences.	74
5.5	Simplified UML-like diagram for AdaptiveUIAdapter	75
5.6	High-level class diagram for Input Adapter Layer data structures.	76
5.7	Adapter WebSocket lifecycle with reconnect backoff.	76
5.8	Real-time adaptation applied after miss-tap event for visually impaired user.	84
5.9	Backend Injection Interface: configuration and live JSON visualization.	85
5.10	Example adaptation list before/after a pointing gesture event.	86
5.11	Example SwiftUI app before/after a miss_tap event.	87
6.1	P1 (motor) miss_tap. The adaptation enlarges the button and border, reducing target difficulty. This aligns with P1’s needs and with the objective scores in Table 6.6 (PAA 50.88%, ERA 57.14%, DCI 0.986).	98
6.2	P5 (visual+motor) slider_miss. The adaptation increases slider size, text contrast and more as seen by the notification in the app. This matches P5’s mixed needs and the strong scores in Table 6.6 (PAA 77.42%, ERA 64.29%, DCI 1.000).	98

List of Tables

3.1	Adaptations mapped to WCAG criteria and intent.	38
4.1	Exact agent settings for the three measured configurations.	63
4.2	Latency and correctness across configurations (sequential backend; <code>user_seq</code> with default profile and growing history).	63
5.1	Adaptation-to-widget mapping (from <code>applyAdaptations</code>).	73
5.2	Adapter endpoints used by <code>adaptive_ui_adapter.dart</code>	76
5.3	MA-SIF agents, focus, allowed actions, and model settings (runtime configurable).	79
5.4	Rule triggers and conservative adaptations (fallback path).	80
5.5	Scope of adaptations in the Flutter implementation.	83
6.1	User profile mapping for evaluation runs.	91
6.2	MA-SIF (balanced) agent configuration.	92
6.3	Mapping of adaptation actions to accessibility categories.	92
6.4	Mapping of error events to acceptable corrective adaptations.	93
6.5	Fitts law proxy: relative reduction in index of difficulty for $s=1.5$	94
6.6	Objective accessibility and design quality metrics by profile.	94
6.7	Overall accessibility targeted actions from the global action distribution.	94
6.8	Policy level WCAG coverage addressed by proposed adaptations.	95
6.9	ERA by event type (higher is better).	95
6.10	Stability across runs (Jaccard similarity of action and target sets).	95
6.11	Counterfactual PAA under profile swaps (mean over all swaps).	96
6.12	Invalid responses by profile (derived from client-side schema check of final payloads).	97
6.13	Rule baseline vs. LLM by profile. Jaccard is set overlap of action names.	97
6.14	Overall feasibility results with MA-SIF (balanced).	99
6.15	Per profile schema validity and latency under MA-SIF (balanced).	99
6.16	Overall configuration comparison (lower latency is better; higher Schema-valid, PAA, ERA, and DCI are better).	100
6.17	Per-profile configuration comparison (p50 latency, Schema-valid, PAA, ERA, and DCI).	100
7.1	Comparison of related adaptive UI systems	104

Chapter 1

Introduction

1.1 Background and Motivation

Whether it's a smartphone app, a website, or a smartwatch, users primarily interact with technology through graphical user interfaces (GUIs). These interfaces rely on visual elements such as icons, buttons, and menus to enable intuitive and efficient interaction, replacing the need for complex text-based commands. Before GUIs became standard, users had to operate computers through command-line interfaces, requiring memorization of commands and technical know-how. This barrier limited computing to experts. GUIs removed that limitation by introducing visual metaphors, opening up digital technology to the general public. Since their introduction in the 1980s, GUIs have revolutionized how humans engage with computers, transforming once-specialized machines into everyday tools. As the critical interface layer between human and machine, GUIs translate complex computational tasks into accessible, visually manageable actions.

Over the decades, GUIs have evolved far beyond their early desktop roots. From the classic WIMP model; windows, icons, menus, and pointers, interfaces have expanded to include touch, voice, gesture, and even spatial interactions in augmented and virtual reality. Today's user interfaces are increasingly multimodal, allowing for richer and more natural interactions. Yet, despite this evolution in interaction modalities, the underlying structure of most GUIs remains largely static. Interfaces are typically designed with a one-size-fits-all approach, offering the same layout, behavior, and visual elements to all users, regardless of their needs or context. This static nature becomes a barrier in today's increasingly diverse and dynamic world, especially for users with accessibility challenges. People with visual impairments, motor difficulties, or those who rely on hands-free interaction methods often struggle with standard interfaces that do not adapt to their specific needs. For example, small buttons can be difficult to tap for users with tremors, while low contrast or lack of screen reader support limits usability for visually impaired users. Moreover, users interact with devices in various environments some noisy, some bright, some on the move, further complicating interaction for those needing alternative input or output modalities. Devices range widely in size and capabilities, from smartwatches and smartphones to desktop monitors and VR headsets, but accessibility features often remain limited or inconsistent across platforms. This gap highlights a key challenge in interface design: how can we create graphical user interfaces that dynamically adapt to individual users' accessibility needs and context, ensuring inclusive and efficient interaction for all?

Adaptive UIs are user interfaces that can adjust themselves dynamically to accommodate the user's accessibility needs and contextual factors. These interfaces are designed to intelligently modify their layout, behavior or visual appearance based on individual user preferences, abilities, and environmental conditions. By responding to multimodal inputs such as touch, voice, gestures, and gaze, adaptive UIs aim to provide a more personalized, inclusive, and efficient user experience, especially for users with diverse accessibility requirements. When we call an interface "intelligent" we mean it can perceive, interpret, and respond to complex user behaviors and contexts autonomously. This level of responsiveness goes beyond simple preset rules, necessitating systems that learn from user interactions and adapt close to or in real-time. Artificial Intelligence (AI), especially advances in machine learning and large language models, offers powerful tools to enable such intelligent adaptation. By processing multimodal inputs, UI context and user data, AI-driven interfaces can dynamically tailor themselves to meet diverse user needs, making accessibility more effective and seamless.

Large Language Models (LLMs) like GPT have revolutionized natural language understanding, processing and generation. Their ability to comprehend context, infer intent, and produce human-like responses makes them highly suitable for interpreting user intent, preferences, and even subtle cues from multimodal inputs such as speech, gestures, or eye movements. In the context of adaptive UIs, LLMs can act as intelligent controllers that translate diverse and complex user interactions into actionable interface adaptations. For example, an LLM can understand a voice command like “make buttons bigger” or interpret hesitation in gestures to trigger UI changes that enhance usability for motor-impaired users. Furthermore, LLMs can dynamically generate or modify UI elements by turning the interface itself into an API enabling on-the-fly customization that is both personalized and context-aware.

Despite growing research and technological capabilities, adaptive user interfaces remain rare in practical use, often due to complexity in implementation and lack of robust frameworks. This gap motivated this thesis, inspired in part by emerging ideas on treating UIs as APIs controlled by intelligent models, an approach highlighted in recent discussions in the AI and HCI communities. Large Language Models (LLMs), in particular, offer a promising interface for driving these intelligent UI adaptations due to their ability to reason over context and user intent. However, current LLMs also come with challenges: latency when used via APIs, potential for hallucinations, a lack of specialization for GUI reasoning or code-generation tasks (for example, rewriting UIs at runtime). These limitations highlight the need for a structured framework that can integrate intelligent models into adaptive UI systems while managing their weaknesses effectively. Rather than seeing LLMs as perfect agents, this thesis treats them as powerful, context-aware assistants that when properly guided and constrained, can significantly improve how interfaces adapt to users’ real-time accessibility needs.

1.2 Problem Statement

Despite increasing attention on accessibility and personalization in user interfaces, most applications remain static and poorly adapted to users’ changing contexts or abilities. While current interfaces can adapt in high-level ways such as adjusting to screen size, orientation or theme, these adaptations are mostly cosmetic. They do not address the deeper challenge of understanding and responding to the user’s intent, cognitive state, or physical limitations in real time.

Designing truly adaptive UIs remains uncommon, in part due to the technical difficulty of modeling dynamic behavior and the lack of robust, general-purpose frameworks. Furthermore, real-world interaction is often messy and unpredictable: small variations in input such as a mistap, delayed reaction, or environmental noise can result in significantly different needs or outcomes. This non-linearity mirrors the principles of chaos theory, where minor initial differences lead to divergent results, making rule-based UI design fragile.

Addressing this requires systems capable of interpreting nuanced, multimodal signals and adapting accordingly, a task well suited to modern AI techniques such as large language models (LLMs). Nevertheless existing UI frameworks rarely integrate such models in a way that supports real-time, intelligent adaptation, leaving a significant gap between the potential of adaptive interfaces and their current practical application.

1.3 Research Objectives

The primary objective of this research is to design, implement, and validate a modular, multimodal AI-driven framework for dynamic user interface (UI) adaptation, with a strong focus on accessibility and personalization. The framework aims to move beyond static or one-size-fits-all designs by enabling real-time, context-aware adaptations that respond to diverse user abilities, preferences, and situational needs. It does so by combining multiple input modalities with hybrid reasoning, rule-based logic for predictable low-latency responses, and large language model (LLM) reasoning for complex or ambiguous cases within a unified architecture.

The overarching goal is to create a system that is both technically extensible and practically deployable across different platforms, while demonstrating tangible accessibility benefits for key user groups, such as motor-impaired, visually impaired, and hands-free users. This will be validated through a proof-of-concept application and simulated evaluation metrics that assess accuracy, effectiveness, and usability impact.

To achieve this goal, the following specific objectives are defined:

1. **Develop a modular, cross-platform architecture:** Design a three-layer framework (Frontend, Input Adapter, and Backend) that can capture, standardize, and process interaction events across Flutter, SwiftUI, and future platforms (e.g., Unity for VR/AR), using a shared JSON-based event and adaptation format.
2. **Implement multimodal input fusion:** Support and combine inputs from touch, keyboard, voice, gestures, and mock or future modalities (e.g., gaze tracking), enabling richer context capture for adaptation decisions.
3. **Integrate Smart Intent Fusion (SIF):** Develop and integrate a hybrid reasoning engine that fuses profile data, interaction history, and live events to generate targeted adaptations, blending deterministic rules with multi-agent LLM reasoning.
4. **Deliver accessibility-focused adaptations:** Provide near real-time UI changes such as enlarging targets, enhancing contrast, switching interaction modes, and adjusting layout or navigation flow, driven by the needs of specific user groups.
5. **Provide a developer-friendly integration path:** Expose the framework as a basic and simple SDK with clear contracts, minimal setup, and reusable patterns, allowing developers to retrofit existing apps or build new ones without major rewrites.
6. **Evaluate system performance and impact:** Assess adaptation accuracy, quality, and latency across different SIF configurations and user profiles through a feasibility study, identifying strengths and areas for future improvement.

1.4 Thesis Structure

This section outlines the organization of the thesis to guide the reader through its content and structure. The thesis is structured as follows:

- **Chapter 1: Introduction**

Introduces the research problem, motivation, and scope of the work. Defines the challenges of static, non-adaptive UIs for accessibility and personalization, outlines the research objectives, and presents the contributions of the thesis. Concludes with an overview of the thesis structure.

- **Chapter 2: Background and Related Work**

This chapter reviews the state-of-the-art in human-computer interaction (HCI), focusing on adaptive user interfaces, multimodal input processing, and (AI-driven) personalization. It discusses existing frameworks for accessibility, user profile built UIs and limitations of static UIs, and the role of large language models (LLMs) in UI adaptation, positioning the proposed framework’s novelty.

- **Chapter 3: System Design and Architecture**

This chapter describes the overall architecture of the framework, including its three-layer design (Frontend Layer, Input Adapter Layer, and Smart Intent Fusion Backend). Explains the role of each layer, the JSON-based event and adaptation contracts, and how the system captures and processes multimodal inputs. It also introduces the Smart Intent Fusion (SIF) architecture and adaptation pipeline.

- **Chapter 4: Smart Intent Fusion (SIF)**

This chapter delves into the specifics of the Smart Intent Fusion (SIF) component, outlining its architecture, the reasoning processes it employs, and how it integrates with the other layers of the framework. It discusses the challenges of fusing diverse input modalities and the strategies employed to ensure accurate and context-aware adaptations as well as dealing with hallucinations from LLMs.

- **Chapter 5: An Adaptive Multimodal GUI Framework using LLMs**

This chapter presents the implementation of a proof-of-concept for the framework, covering the Flutter-based “Adaptive Smart Home Controller” UI with touch, and simulated voice/gesture inputs, the input adapter layer’s JSON contract, and the SIF backend. A backend Flutter-based interface is also introduced for developer insights and debugging.

- **Chapter 6: Feasibility Study**

This chapter discusses the evaluation setup, scenarios, results, and reflects on the implications and limitations of the findings.

- **Chapter 7: Discussion and Future Work**

This chapter discusses the broader implications for accessibility and HCI, key findings, and comparison with related work. Addresses limitations and proposes future work, including on-device reasoning, integrated visual UI understanding, federated learning, and expanding modality support.

- **Chapter 8: Conclusion**

This chapter summarises the thesis contributions and lessons learned from the development process. Reflects on the evaluation results, revisits the research objectives, and outlines how this work can inform future adaptive UI systems.

Chapter 2

Related Work

2.1 Multimodal AI in User Interfaces

User interfaces have evolved far beyond simple point-and-click paradigms. Modern systems increasingly incorporate **multimodal inputs**, integrating different sensing and communication modalities such as voice or speech commands, pen and touch interaction, eye tracking, and full-body gesture tracking. These technologies aim to create more natural, flexible, and inclusive experiences [36]. Multimodal interfaces aim to mirror human communication, which rarely relies on a single channel, thereby expanding the possibilities for interaction and enabling more robust and adaptive user experiences.

A fundamental advantage of multimodal interfaces is their potential to improve accessibility and inclusivity. For example, voice commands can be crucial for users with motor impairments, while gaze-based control may assist users who cannot use their hands for input. Meanwhile, gestural input like in VR settings can support contactless control, useful in medical environments, during physical rehabilitation exercises or even astronaut training. Furthermore, the combination of modalities can reduce error rates and cognitive load by providing redundant and complementary channels of communication, improving user satisfaction and performance [36].

Historically, the field of multimodal interaction gained attention through early research prototypes combining speech and pen input, which demonstrated that parallel and redundant channels lead to more fluid interaction [36, 37]. In recent years, rapid advances in machine learning, especially deep learning, have further enabled real-time recognition of speech, gestures, and gaze on consumer-grade hardware [27, 9].

2.1.1 Pointing Devices and Touch Interfaces

Before the rise of multimodal and intelligent input modalities, most user interfaces primarily relied on pointing devices such as the mouse, trackball, and touchpad. Fitts' law [16] and subsequent pointing models formed the foundation for optimizing target sizes and layouts to reduce pointing time and error rates. Touchscreens expanded on these paradigms, introducing direct manipulation, gesture-based scrolling and multi-touch interactions. Touch interfaces, now ubiquitous in smartphones and tablets, benefit from intuitive mappings between finger movements and on-screen actions but also face challenges related to occlusion, precision, and physical fatigue [44].

Recent research has introduced adaptive techniques, such as "sticky" or magnetic cursor effects, which are already in use today as normal or accessibility features on for example Apple devices. Which help guide the pointer toward interactive elements, reducing effort and error, particularly for users with motor impairments [10]. These principles guide contemporary accessibility enhancements and encourage advancements in multimodal environments. In this thesis, inspiration is drawn from those adaptive cursor behaviors to inform dynamic UI element resizing and magnetic effects in multimodal contexts, bridging classical pointing interactions with "smart" AI-driven adaptations.

2.1.2 Voice-Driven Interfaces

Voice interaction has become mainstream through virtual assistants like Amazon Alexa, Google Assistant, and Apple’s Siri. These systems utilize automatic speech recognition (ASR) and natural language processing (NLP) to interpret user commands and execute them on the device [21]. Voice input offers hands-free accessibility benefits, particularly useful for users with mobility impairments or in multitasking scenarios. However, it also presents challenges such as interpreting unclear or ambiguous speech, addressing privacy concerns, and maintaining reliability in noisy environments [33]. Combining voice with complementary modalities such as gestures or gaze tracking has been shown to improve disambiguation and user confidence [32]. Hybrid systems support error correction, implicit confirmation, and increased interaction bandwidth [37], thereby enhancing overall usability.

2.1.3 Gesture and Gaze Integration

Gesture-based interactions enable users to control interfaces through hand and body movements, providing expressive and intuitive command options. Advances in computer vision techniques and frameworks such as MediaPipe Hands and OpenPose have made real-time gesture recognition feasible on consumer-grade hardware including desktops and mobile devices [27]. Eye tracking is another important modality, capturing user attention and intention. It has been used to implement gaze-contingent interfaces where UI elements respond dynamically to where users are looking, reducing physical effort and increasing interaction speed [13]. The integration of gesture and gaze input creates a powerful multimodal system, enabling users to point, select, and confirm actions more naturally. Studies indicate that combining these modalities improves error tolerance and enhances overall user experience, especially in contexts where precise manual input is difficult or limited, such as virtual reality environments [35, 7].

One of the most influential early systems that explored the practical use of gaze in everyday computing contexts is the GUIDe project by Kumar and Winograd (2007) [23], which demonstrated how gaze could be used as a lightweight augmentation of traditional input methods rather than a full replacement. Their EyePoint technique introduced a refined two-step “look-press-look-release” interaction loop that allowed for highly accurate gaze-based pointing by combining magnification with gaze refinement. Unlike previous systems that relied solely on dwell-based activation, GUIDe’s use of keyboard-assisted targeting reduced false activations (the “Midas Touch” problem) and allowed for precise user control.

They extended this principle to application switching (EyeExposé) and adaptive scrolling (EyeScroll), showing that gaze input could become a viable interaction modality for able-bodied users if the design balanced gaze with explicit control channels like hotkeys or contextual awareness. Importantly, EyeScroll demonstrated adaptive scrolling speeds based on live gaze-driven reading behavior, which is a strong precedent for personalization and multimodal input adaptation. These contributions illustrate how gaze can serve as both a passive signal and an active intent channel, especially when fused with complementary modalities like touch or speech [24].

2.2 Adaptive GUIs Across Modalities and Platforms

As user interfaces evolve to support an increasing variety of input modalities and devices, adaptive graphical user interfaces (GUIs) have gained significant attention. Adaptive GUIs are designed to modify their layout, behavior, and appearance dynamically in response to contextual information, user preferences, or real-time interaction patterns. This adaptivity aims to enhance usability, accessibility, and personalization across a wide range of platforms, from traditional desktop systems to mobile devices and immersive virtual reality (VR) environments.

Adaptation can take many forms. For example, a GUI might increase button sizes for users with motor impairments, change color schemes to improve contrast for users with visual impairments, or rearrange content based on the user’s task context or dominant hand [18]. On desktop and mobile devices for example, adaptivity has often focused on supporting accessibility and optimizing layouts for different screen sizes or resolutions. Responsive design frameworks, which automatically adjust content and element placement across devices, represent one form of early structural adaptivity. More advanced adaptive systems might learn user habits over time and present solutions based on the users context by for example integrating large language models (LLMs) and advanced AI reasoning.

2.2.1 Accessibility-Focused Adaptive GUIs (Desktop & Mobile)

Adaptive GUIs on conventional platforms personalize layout, component size, and interaction mechanics to user ability and context. A canonical example is SUPPLE, which models devices, tasks, preferences, and motor abilities to automatically generate alternative UIs e.g., larger widgets, altered control density, or layout re-flow for users with limited range of motion [19]. Complementing such generation, Ability-Based Design argues the system should adapt to people’s abilities (not vice-versa), yielding techniques like dynamic gain adaptation for pointing, larger targets, and stabilization aids that improve accuracy without burdening the user [45]. In everyday menus, empirical work shows that adaptable or carefully designed adaptive structures can speed selection compared to static menus when they preserve predictability and user control [14, 40]. More nuanced approaches such as ephemeral adaptation highlight predicted items while maintaining spatial consistency to reduce visual search cost [15]. Together, these strands illustrate how adaptive GUIs can enlarge controls, re-order or highlight options, and subtly modulate input mechanics to deliver accessibility and personalization on desktop and mobile.

2.2.2 Design Challenges for Adaptive GUIs Across Modalities

Designing effective adaptive GUIs faces well-documented pitfalls. Classic guidance warns that opaque or unpredictable adaptations erode trust; users need scrutability and control over what changes and why [22, 14]. Visual stability matters: aggressive re-ordering can harm spatial memory, hence techniques that preserve layout while cueing salience (e.g., ephemeral onset) are preferred [15]. Input-side adaptations must be unobtrusive for non-target users yet beneficial for those who need them for instance, target-agnostic cursor gain adjustments improve motor-impaired pointing while remaining neutral for others [46]. Finally, evaluating accessibility-oriented adaptations requires careful methodology and representative sampling; recent surveys highlight gaps in reporting, participant diversity, and measures for motor-impairment studies, issues framework’s standardized event contracts and profiles can help address [38]. Overall, successful systems balance transparency, user agency, visual consistency, and rigorous evaluation to ensure adaptations help rather than hinder.

2.3 Classical Adaptive UI Techniques

Before the widespread use of artificial intelligence and machine learning in dynamic interfaces, adaptive graphical user interfaces (GUIs) primarily relied on deterministic, rule-based strategies. These classical approaches, while limited in personalization and flexibility, were instrumental in shaping the early landscape of user interface adaptation [34, 8]. They focused on fixed logic, device-specific configurations, and contextual parameters defined explicitly by designers or developers. Although static by today’s standards, these techniques addressed fundamental issues of usability, accessibility, and device heterogeneity.

One of the most foundational classical techniques is responsive layout design, which allows a user interface to adjust its layout and elements based on screen size, orientation, and device type. This approach is especially prominent in web and mobile design, where CSS media queries and related standards dynamically reposition and resize UI components to preserve usability across different platforms [30, 25]. Alongside responsive design, constraint-based layout systems offered deterministic ways to maintain relationships among UI elements, enabling consistent resizing and alignment without learning from user behavior [3]. While these methods improve accessibility and device compatibility, they do not adapt to individual user behavior or preferences, the adaptations are driven solely by environmental conditions like screen dimensions or device class.

Another key classical technique is context-aware UI adaptation, where the interface adjusts based on predefined environmental or situational variables such as location, time of day, network connectivity, or battery level. Seminal frameworks and surveys formalized what “context” is and how it can be used to trigger if-then rules and proactive adjustments in applications [11, 4]. For example, a mobile app might switch to a dark theme automatically at night or reduce update frequency on a metered connection. These adaptations are typically rule-based rather than probabilistic, improving relevance and efficiency but lacking the ability to learn from user interaction patterns over time.

Additionally, role-based adaptations were widely used in enterprise and healthcare applications to simplify complex systems for distinct user groups. Interfaces were tailored to the responsibilities and permissions of roles (e.g., administrator, clinician, student), exposing only the necessary features and layouts for

each group and thereby reducing cognitive load [1, 26]. Such role-driven tailoring was predetermined by organizational needs rather than inferred intent, exemplifying classical, deterministic adaptation.

2.3.1 Responsive Layouts

Responsive design emerged as a critical approach to address the rapid increase of different screen sizes and devices, particularly with the rise of smartphones and tablets. Instead of creating separate fixed interfaces for each device, responsive layouts allow a single design to automatically adjust its elements to fit various screen dimensions and orientations based on core techniques like flexible grids, media queries, and adaptive content strategies. Those enable dynamic resizing and rearrangement of UI elements [31]. In mobile development, declarative UI frameworks such as Flutter and SwiftUI adopt similar principles, using constraints and reactive layouts that adapt hierarchically to screen changes. While originally developed to support multiple screen sizes, responsive layouts also improve accessibility by allowing users to zoom or scale interface components without breaking layout integrity. Despite their flexibility, traditional responsive designs are typically static in behavior they do not react to real-time user behavior or contextual signals beyond the device dimensions.

2.3.2 Context-Aware Design

Context-aware interfaces extend the notion of adaptation by responding to more complex environmental and user-specific factors. Originating in ubiquitous and pervasive computing research, context-aware design involves sensing parameters such as location, time of day, user activity, nearby devices, or even physiological signals to modify interface behavior [39, 11]. For example, a context-aware mobile application might switch to a dark theme automatically at night or suggest different UI shortcuts when the user is driving versus sitting at a desk. In smart home systems, interfaces may change based on proximity sensors or room occupancy data. Some adaptive systems also integrate user profiles and long-term behavior patterns. For example, the SUPPLE system automatically generates personalized interfaces optimized for a user’s specific motor abilities and preferences by solving optimization problems over interface layouts with the effect of facilitating faster access and lower error [18]. Such rule-based or optimization-based approaches set the stage for later, more sophisticated AI-driven adaptations by highlighting the importance of context and user modeling.

While classical techniques like responsive layouts and context-aware design lack the semantic understanding and reasoning capabilities of modern AI systems, they provide robustness and predictability. They form a crucial baseline against which the benefits of AI-enhanced adaptations can be evaluated. Furthermore, they remain relevant in many production systems due to their lower computational cost and easier predictability for developers and designers.

2.4 Programmable UIs

Traditionally, user interfaces have been treated as static artifacts, closely tied to fixed layouts and hard-coded interaction patterns. However, recent advances in computational understanding of user interfaces propose a shift towards making UIs “programmable” that is, representing them as dynamic, semantically rich structures that can be analyzed, modified, and generated automatically. This paradigm enables interfaces to become more flexible and adaptable, allowing developers and even AI systems to reason about and transform UI elements in response to changing contexts, user needs, or platform constraints. Such a programmable layer is essential for realizing fully adaptive, AI-driven interfaces that can personalize experiences and improve accessibility in real time.

2.4.1 Reflow

A notable example of pixel-based UI adaptation is Reflow, introduced in Chapter 7 of Wu’s dissertation on computational understanding of user interfaces [47]. Reflow proposes a system that automatically refines and optimizes touch interactions in mobile applications using only pixel-level information, without requiring access to the underlying application code or view hierarchy like a widget tree. The system operates by first detecting UI elements from a screenshot using machine learning-based pixel analysis. It then refines the layout based on a user-specific spatial difficulty map, which identifies the difficult-to-access areas of the screen, derived from calibration tasks that capture individual motor abilities and

preferences. Finally, Reflow re-renders a modified version of the UI with optimized element positions and sizes, ensuring that critical interactive components are easier to reach and select.

Reflow exemplifies how computational understanding of UIs at the pixel level can enable automated personalization and accessibility enhancements even in closed-source or inflexible applications. It illustrates a promising direction for future adaptive systems that aim to provide user-specific improvements without requiring cooperation from original app developers or extensive code modifications. This aligns directly with the goals of AI-driven UI frameworks focused on dynamic, user-centered adaptation. Furthermore, the user study showed an average increase of 9% in interaction speed as well as improved interaction speeds by up to 17% [47].

2.4.2 UICoder

A further advancement in making user interfaces programmable is presented through UICoder, as described in Chapter 8 of Wu’s dissertation on computational understanding of user interfaces [47]. While earlier systems such as Reflow focused on pixel-level adaptations of existing applications, UICoder tackles the challenge of generating high-quality UI code directly from textual descriptions using large language models (LLMs). UICoder addresses two major limitations of previous code-generation approaches for UIs: the scarcity of high-quality, self-contained UI code in existing datasets and the difficulty LLMs face in incorporating visual or spatial feedback into their training. Rather than relying on manually curated or external datasets, UICoder introduces an automated method to iteratively generate, filter, and refine synthetic UI code datasets.

The system begins by prompting an existing LLM to generate large collections of UI code samples (specifically SwiftUI, Apple’s coding framework) from textual descriptions. These outputs are aggressively filtered and scored using compilers (to ensure syntactic correctness) and vision-language models (to assess visual relevance), producing a refined dataset for further finetuning. By iteratively repeating this cycle generation, filtering, and finetuning UICoder progressively learns to produce syntactically correct and visually coherent UI code. Derived from the StarCoder family (which lacked extensive SwiftUI training data), UICoder ultimately generated around one million SwiftUI programs over five iterations. Despite these constraints, it significantly outperformed all open-source baselines and approached the performance of larger proprietary models in both automated and human evaluations. By leveraging automated feedback instead of costly human annotations, UICoder demonstrates a scalable approach to training LLMs for UI code generation. This method shows how generative models combined with programmatic refinement loops can enable on-demand creation of high-quality, personalized UIs directly supporting visions of adaptive and dynamically generated interfaces.

UICoder exemplifies the shift from static UI codebases toward dynamic, generatively defined interfaces. In the context of AI-driven adaptive GUIs, such a capability enables systems to not only adjust existing layouts but also generate entirely new interface code on demand, tailored to specific user preferences or device contexts. This directly supports the vision of self-adaptive and user-personalized UI frameworks proposed in this thesis.

2.4.3 User Interface Adaptation using Reinforcement Learning

Recent advances in programmable and adaptive user interfaces have explored the integration of machine learning techniques particularly Reinforcement Learning (RL) to support real-time personalization and dynamic behavior. A notable example is the doctoral work by Gaspar-Figueiredo (2023), which proposes a UI adaptation framework that leverages RL in conjunction with physiological data to enhance user experience (UX) [20].

In this framework, RL is used to determine optimal UI adaptations (e.g., layout or content adjustments) by continuously interacting with the user and optimizing long-term reward signals. What sets this work apart is the use of physiological signals such as eye tracking, EEG, or other biosignals as objective measures of user response. This addresses a key limitation of traditional evaluation techniques that rely on subjective self-reporting, which can introduce bias or fail to capture moment-to-moment changes in experience.

The system learns from users’ physiological reactions and interaction behaviors to determine which adaptations are effective in improving usability and engagement. Over time, the interface becomes more personalized and context-aware, reacting not only to explicit input but also to subtle cues from the

user’s cognitive or emotional state. This approach exemplifies a new direction in programmable UIs: systems that are not only defined by abstract models (e.g., UIML or USiXML) but are also capable of learning and evolving through interaction. By combining adaptive logic, ML-driven decision-making, and, biosignal-based feedback, such frameworks could serve as the basis for next-generation interfaces especially in applications involving accessibility, cognitive load, or health and wellness.

2.5 User Profile built UIs

The increasing diversity of devices and user contexts has driven research into creating adaptable user interfaces that are both device-independent and user-centered. A significant challenge arises in balancing generalization so an interface works across multiple devices and personalization so it aligns with the unique preferences and capabilities of individual users.

To address this, Luyten et al. [29] proposed a framework that combines high-level XML-based user interface description languages (UIDLs), particularly UIML, with MPEG-21 Part 7 (Digital Item Adaptation) user profiles. UIML (User Interface Markup Language) allows designers to abstract the UI from concrete implementations, enabling interfaces to be rendered differently depending on device constraints. For example, a single abstract element like “choice from a range” can map to different widgets (slider, list, or text input) depending on the target platform. In this approach, MPEG-21 user profiles capture individual user preferences and requirements (e.g., accessibility needs, preferred interaction styles), which can then dynamically guide the adaptation of the UI described in UIML. This enables the generation of multi-device, personalized interfaces that are both broadly deployable and tailored to specific user needs.

An implemented prototype demonstrated how combining UIML and MPEG-21-based user profiles allows for seamless adaptation of UI layouts and interactions while minimizing design effort. By leveraging abstract UI definitions and structured user profiles, this method enhances both accessibility and usability across diverse platforms, making interfaces more “granny-proof” and inclusive. This user-centric adaptation model supports the vision of highly personalized digital experiences without sacrificing cross-device compatibility, contributing an important step toward universal, accessible, and adaptable user interfaces.

Early profile-based UI adaptation systems were often limited by the rigidity of XML-based markup languages and the static nature of their user models. Profiles were typically loaded once and assumed relatively stable user needs, which limited the ability to adapt interfaces dynamically over time or in real-time scenarios.

However, these systems laid foundational groundwork for modern AI-augmented adaptation layers. For instance, where MPEG-21 profiles might have been manually filled or gathered via forms, today’s systems can infer similar parameters from behavioral data using machine learning. Similarly, abstraction principles from UIML still resonate in declarative UI frameworks like Flutter, SwiftUI, or React Native, which separate layout logic from presentation and enable device-responsive rendering.

These classical approaches remain valuable not only for historical understanding but also as a reminder of the importance of modularity, abstraction, and structured user modeling in UI design principles that continue to influence the development of multimodal, intelligent frontends today.

2.5.1 XML-Based Runtime UI Systems

One early precursor to programmable and adaptive UI systems was the use of XML-based runtime user interface description languages, especially for resource-constrained mobile and embedded systems. A notable example is the work by Luyten and Coninx (2001), who proposed a method that leverages XML for describing UI components and Java for rendering them on mobile devices such as PDAs [28]. Their system allowed interfaces to be dynamically serialized, transmitted, and rendered on client devices based on contextual variables such as the user’s role, device capabilities, and preferences. For example, the interface for controlling a projector would differ depending on whether the user is a professor or a technician, enabling personalized task-specific controls.

This work introduced several important ideas that prefigure modern programmable and adaptive UIs. First, it treated UIs as dynamic data rather than static views, enabling runtime generation and adaptation. Second, it used constraint-based filtering mechanisms to tailor UI components based on context.

Third, it demonstrated that a separation between the UI’s description (in XML) and execution (in Java) could support modular, reusable interface logic.

Although this approach predates current AI-powered frameworks, it exemplifies early efforts toward what is now called adaptive or context-aware UI. Its emphasis on platform-independence, modularity, and runtime flexibility directly anticipates features found in modern frameworks like Flutter, React, and Unity UI especially when extended with AI-driven mediation and multimodal inputs.

2.6 Multimodal Fusion and Input Event Modeling

As user interfaces become increasingly multimodal, systems must handle diverse streams of input signals in a coherent way. Multimodal fusion refers to the process of integrating these different input modalities such as voice, touch, gaze, and gestures into unified semantic commands. This enables interfaces to interpret combined user inputs more naturally and contextually, mirroring how humans often combine speech and gestures in daily communication.

Closely related, input event modeling abstracts low-level raw data (like finger coordinates, gaze points, or audio waveforms) into higher-level representations of user intent (such as “select”, “scroll”, or “confirm”). This abstraction layer serves as an essential bridge between noisy sensor data and actionable interface responses. By modeling input events at different levels of granularity from raw physical movements to semantic-level intents systems can reason about user behavior, handle ambiguities, and provide more robust feedback.

For example, in a health application, simultaneous voice and gaze inputs might be fused to allow a patient to say “yes” while looking at a confirm button, providing redundancy that improves reliability for users with motor or speech impairments. Similarly, in VR or AR, interpreting combined hand gestures and head movements enables more precise object manipulation and scene navigation, which would be challenging with single-modality input alone.

2.6.1 Fusion Architectures

Several architectures have been proposed to integrate multimodal inputs effectively [2]. Early fusion approaches combine raw input data at a low level, for example merging voice and gesture signals before any individual interpretation occurs. While this can enable richer context awareness, it often requires precise synchronization and robust signal alignment, which can be technically challenging. On the other hand, late fusion architectures process each modality independently to obtain intermediate interpretations (such as recognized words or detected hand poses), and then merge these at a semantic level. This approach tends to be more modular and easier to maintain, as each input channel can be improved or swapped without affecting the others.

Hybrid fusion combines elements of both early and late strategies, allowing certain low-level signals to be shared while keeping higher-level interpretation pipelines separate [5, 36]. Choosing an appropriate fusion architecture is critical for ensuring fluid and error-tolerant interactions. For example, in an adaptive health app, fusing gaze and speech can enable users with motor impairments to confirm commands more easily. Similarly, in AR/VR settings, combining hand tracking with eye gaze supports natural object selection and manipulation.

2.6.2 Event Abstraction Models

Once input signals are fused, systems must convert them into abstracted events that represent user intent rather than raw movements or spoken words. Event abstraction models define a hierarchy of events, from primitive input (e.g., “finger tap at position x,y”) to higher-level semantic commands (e.g., “confirm selection”, “scroll left”, or “open menu”). This abstraction not only simplifies the logic needed to respond to different input combinations but also improves the system’s adaptability across platforms and devices. By mapping diverse physical actions to a common set of abstract commands, a single system can support a broad range of user abilities and contexts.

Recent work in multimodal interaction design has emphasized the importance of flexible and extensible event models that can incorporate new modalities without extensive re-engineering [6]. Moreover, integrating AI-driven models, such as large language models or gesture classifiers, can further enrich the abstraction process by inferring user intentions from subtle or ambiguous input cues. Together, fusion

architectures and event abstraction models form a foundation for building robust, adaptive, and inclusive multimodal user interfaces. As systems continue to evolve, these techniques will play a crucial role in creating personalized and accessible experiences across devices and environments.

A notable example of semantic event abstraction in real-world interfaces is the work of Dixon et al. [12], who implemented a general-purpose, target-aware pointing enhancement based on the Bubble Cursor. Their system uses pixel-level reverse engineering (via Prefab) to identify interface components and overlay interaction semantics onto them, allowing the system to interpret raw pointer movement as intent to interact with semantically meaningful targets even when underlying applications do not expose accessibility metadata. This layered approach of separating visual identification from interaction intent closely mirrors the goals of input event modeling in multimodal systems, especially in contexts where event boundaries are ambiguous or where UI elements are rendered outside standard toolkits.

2.7 LLMs as UI Controllers

Recent advances in large language models (LLMs) have opened new opportunities for bridging natural language and user interfaces. Traditionally, UI control has relied on explicitly designed event handlers and rigid APIs, requiring precise user input or structured interaction patterns. LLMs, by contrast, enable more flexible, high-level interactions that resemble natural human communication, which is especially promising for non-expert users or contexts where accessibility is crucial.

By leveraging LLMs’ powerful capabilities in understanding and generating natural language, it becomes possible to control user interfaces using text or voice instructions without needing predefined UI-specific commands. For example, a user could ask an application to “show me my upcoming appointments and move the next one to next week,” and an LLM-based controller can parse this instruction, map it to the appropriate UI actions, and execute them seamlessly.

This paradigm allows user interfaces to function more like intelligent agents rather than static interaction surfaces, reducing cognitive load and lowering the technical barrier for interaction. Furthermore, LLM-driven UI controllers can adapt to user-specific phrasing and preferences over time, learning from previous interactions to provide more personalized and efficient support.

2.7.1 Turning UIs into APIs

One promising approach to LLM-driven interfaces is conceptualizing user interfaces as implicit APIs. Instead of interacting with the UI through direct manipulation (e.g., clicking buttons, dragging elements), the UI’s functionalities are abstracted into callable actions through an adapter layer that can be invoked through language. This “UI-as-API” perspective treats every interactive element and functionality as an endpoint or command that can be described and triggered textually. As demonstrated in recent research on program synthesis and UI automation like the UICoder referenced earlier [47, 43], LLMs can be trained to translate high-level instructions into structured API calls or internal code representations. This enables a two-layered interaction model: the LLM interprets free-form user instructions and maps them onto the UI’s abstracted actions, which are then executed to update the visible interface.

A significant advantage of this model is that it can help unify interaction modalities. Whether a command is given via voice, text, or even gesture-based language input, it is ultimately funneled into the same set of abstracted API calls. This makes interfaces more robust to modality switching and enhances accessibility. Moreover, turning UIs into APIs facilitates automation and integration with external services, enabling systems to not only react to individual user commands but also orchestrate multi-step tasks and workflows automatically. While challenges remain such as handling ambiguous instructions or ensuring robust error handling and LLM hallucinations, this direction shows strong potential for creating more adaptive, intelligent, and user-centered interfaces.

2.7.2 Agents

The rise of LLM-based agents represents a powerful evolution of user interface control beyond simple command mapping. Unlike static UI controllers, agents leverage LLMs to autonomously interpret, plan, and execute sequences of actions on behalf of users, effectively turning the interface into an intelligent collaborator rather than a passive tool. These agents can reason about user goals, maintain conversational context, and dynamically choose the best sequence of interface actions to fulfill complex or vague

instructions. For example, an agent might respond to "Help me prepare for my upcoming trip" by checking the user's calendar, suggesting packing lists, booking transportation, and even setting reminders all without requiring the user to explicitly navigate each step.

Recent systems such as OpenAI's GPT-based function calling, AutoGPT, or tools like Microsoft's Copilot and Google's Duet illustrate how agents can operate over complex applications, combining high-level reasoning with programmatic UI actions. In research, approaches like ReAct (Reasoning and Acting) frameworks show how agents can chain thought processes ("chain-of-thought") and API-level actions in iterative loops, enabling them to verify outcomes and adapt their behavior [48]. Moreover, LLM-based agents can incorporate user feedback continuously to refine their behavior, creating highly personalized assistants that align closely with individual preferences and working styles.

To further enhance their adaptability, modern agents increasingly integrate environment modeling and multimodal perception, enabling them to not only parse language but also interpret visual interfaces, sensor data, and user gestures. Frameworks like SeeAct [49] and ViperGPT [41] have demonstrated how combining vision-language models with action planning allows agents to operate UIs from visual input alone clicking buttons, reading menus, or navigating unfamiliar applications much like a human would. This opens doors to agent-driven interfaces for accessibility scenarios (e.g., voice or gaze-controlled UIs), remote control of complex software, or even autonomous use of standard desktop/web applications.

2.8 Health and Accessibility Applications

Health and accessibility applications represent some of the most impactful and socially significant domains for adaptive and intelligent user interfaces. As populations age and chronic health conditions become more prevalent, digital health tools are increasingly critical for supporting independence, self-management, and personalized care. Similarly, accessibility-focused interfaces help reduce barriers for users with diverse abilities, ensuring equitable access to technology. Modern health applications leverage multimodal inputs and adaptive interfaces to create more engaging and supportive experiences. For example, apps for physical rehabilitation frequently combine motion tracking (via cameras or wearable sensors) with adaptive visual feedback to guide patients through exercises safely and effectively. Examples include tools like Kaia Health [42] for musculoskeletal therapy and Reflexion for cognitive and physical rehabilitation, which adjust exercise difficulty and feedback based on real-time performance data.

In the mental health domain, conversational agents and virtual coaches are becoming increasingly popular. Applications like Woebot or Wysa utilize natural language processing to provide immediate, conversational mental health support. By continuously adapting their conversational style and recommendations to the user's emotional state and progress, these systems illustrate the power of dynamic, user-centered design. Research has shown that such adaptive, agent-based interactions can improve adherence and patient outcomes [17]. Accessibility applications also demonstrate the importance of personalized, context-aware interfaces. For users with visual impairments, screen readers and AI-powered image descriptions enable rich content access. Gaze-based or switch-based interaction systems empower users with severe motor disabilities to control complex interfaces using minimal input. Projects like Apple's VoiceOver and Microsoft's Seeing AI show how integrating multimodal AI into accessibility solutions can drastically improve daily usability and independence.

Furthermore, combining health and accessibility perspectives opens opportunities for fully personalized assistive systems. For instance, an intelligent multimodal interface could dynamically adjust text sizes and contrast for a user with low vision while also simplifying navigation for cognitive accessibility and providing voice guidance tailored to the user's speech patterns or preferences. Looking toward future developments, advances in virtual reality (VR) and immersive technologies can further enhance these assistive systems. In VR-based rehabilitation, for example, adaptive multimodal interfaces can create engaging, safe, and highly individualized therapy environments, dynamically adjusting exercise difficulty, feedback modalities, and visual cues to match each patient's needs and progress. Such systems show promise for applications ranging from motor skill recovery after stroke to anxiety reduction and pain management.

Chapter 3

System Design and Architecture

3.1 Introduction to System Design

The adaptive multimodal GUI framework developed in this thesis is designed to deliver personalized, accessibility-focused interface adaptations in real time. It builds on the idea that different users have different needs, and that these needs can change depending on context, device, and input modality. By combining multiple input channels such as touch, voice, and gestures with AI-driven reasoning, the framework can adapt interfaces in a way that is both responsive and context-aware.

The architecture follows a three-layer design. The frontend layer renders the user interface, captures user interactions, and applies adaptations as instructed by the backend. The input adapter layer standardizes events across modalities into a common JSON schema, ensuring consistent processing regardless of their origin. The backend layer processes these events using Smart Intent Fusion (SIF), which combines rule-based logic with multi-agent LLM reasoning to generate targeted adaptations. This layered approach was chosen to ensure modularity, scalability, and extensibility for future modalities. By separating concerns across layers, the framework can easily integrate new input methods or adapt to different platforms without significant rework. Each layer can evolve independently, allowing for targeted improvements and innovations.

This chapter presents the framework at a conceptual level, focusing on its architecture, data flow, goals and key design principles. Detailed implementation aspects, such as widget properties, SIF, API routes, or database queries, are deferred to Chapters 4 and 5.

3.2 Overview of the System Architecture

The framework follows a three-layer architecture that was chosen and designed to be modular, scalable, and adaptable to diverse accessibility needs. Each layer has a clearly defined role and communicates with the others through a common JSON-based event and adaptation format. This separation of responsibilities makes it possible to extend or replace individual components without disrupting the rest of the system, and ensures that adaptations can be applied consistently across platforms and modalities.

At the top, the Frontend Layer is responsible for rendering the user interface, capturing interactions, and applying adaptations received from the backend. The input adapter layer sits between the frontend and backend, converting raw interaction data from multiple modalities into the shared JSON schema so that all events are processed in the same way. The backend layer implements the Smart Intent Fusion (SIF) process, combining event data, user profiles, and recent interaction history to produce personalized adaptation actions. The backend can rely on both deterministic rules and multi-agent LLM reasoning, described in detail in Chapter 4. The system is built around a feedback loop. Each user interaction is captured, processed, and logged along with the resulting adaptation. This log contributes to the user's profile and informs future adaptation decisions, allowing the interface to become progressively more personalized over time.

To illustrate this flow in practice, consider a motor-impaired user attempting to press a “Lock” button. If the tap misses its target, the event is sent through the adapter, recognised as a miss-tap, and forwarded

to the backend. SIF responds by generating an adaptation that increases the size and spacing of the button. The frontend receives this instruction and animates the change in real time. For a hands-free user, the process might begin with a spoken command to “turn on the lamp.” The voice input is captured, standardised, and sent to the backend, which returns instructions to switch the UI into voice mode and activate the lamp control.

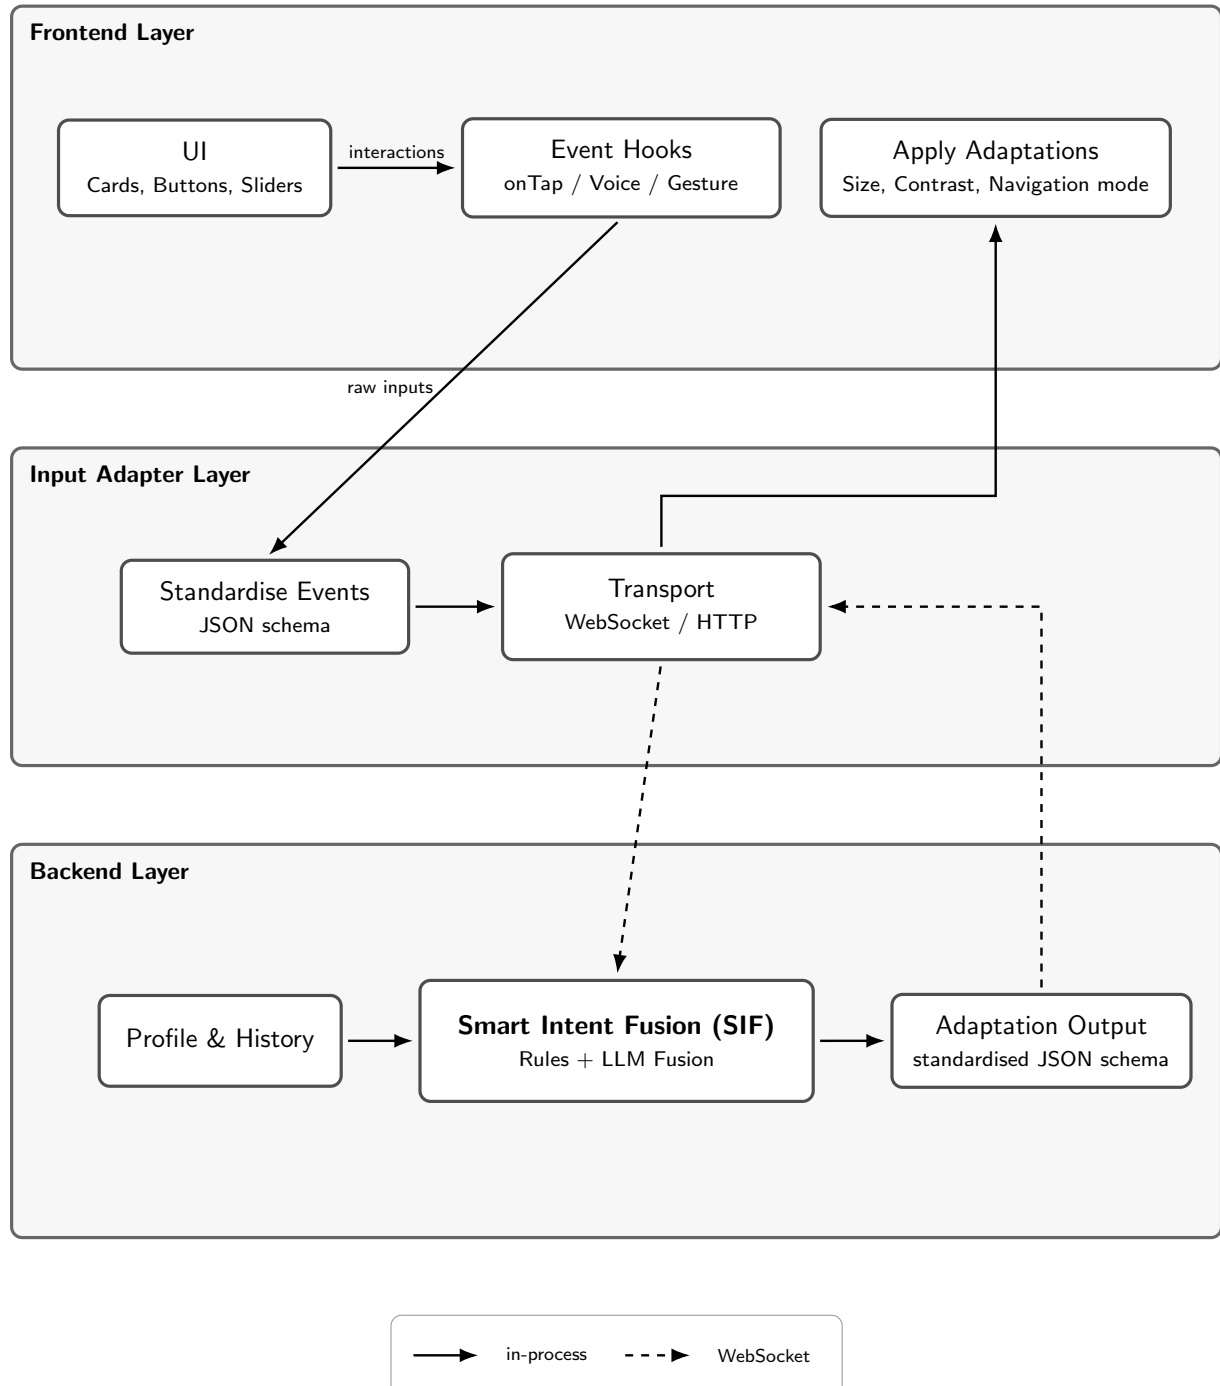


Figure 3.1: High-level architecture flow: frontend events are standardised and sent to the backend, where reasoning fuses user context and generates adaptation outputs for the UI.

The high-level architecture diagram (Figure 3.1) illustrates the event flow from the Frontend Layer, where user interactions are captured, through the input adapter layer that standardizes events into a JSON schema, and onward to the backend layer for Smart Intent Fusion (SIF) reasoning. Adaptation outputs are sent back to the frontend via WebSocket for real-time UI updates, while HTTP endpoints handle profile management and non-interactive operations. The diagram highlights clear separation of responsibilities, modular extensibility for new modalities or reasoning agents, and platform-agnostic communication, ensuring that the event-adaptation pipeline remains robust and easily extendable across domains.

Key Design Principles:

- **Modularity:** Layers can be swapped or upgraded independently.
- **Scalability:** Async processing and MongoDB indexing handle high interaction volumes.
- **Generalizability:** Platform-agnostic design enables deployment in domains from smart homes to healthcare.
- **Accessibility Focus:** All adaptations are guided by WCAG 2.1 and target motor-impaired, visually impaired, and hands-free users.

The most challenging aspect of this design is ensuring seamless communication between layers while maintaining low latency for real-time adaptations, as well as carefully integrating LLMs as a central component for processing and understanding user intents. This will bring challenges on its own such as hallucinations or misinterpretations of user input. Making sure the LLM accurately captures user intent and context is crucial for effective adaptations. Prompt engineering and careful design like fallback options will be essential to mitigate these issues.

3.2.1 Accessibility Focus and Target User Groups

While the framework is general enough to support a wide range of adaptive UI scenarios, its design in this thesis is intentionally centered on three key user groups: motor-impaired, visually impaired, and hands-free users. These groups were chosen because they represent distinct accessibility challenges that can be addressed effectively through multimodal interaction and real-time adaptation.

Motor-Impaired Users Motor impairments can make precise pointing, dragging, or rapid tapping difficult. For these users, repeated failed attempts at interacting with a UI element (e.g., a small button) can lead to frustration and reduced task completion rates. The framework addresses this by detecting such patterns and applying adaptations like enlarging target sizes, increasing hitbox areas, or enabling alternative input modes such as voice or gesture. These adaptations can be applied temporarily (for the current session) or persistently (as part of the user’s profile), depending on the severity and frequency of the interaction difficulties.

Visually Impaired Users Users with partial vision loss or low contrast sensitivity often struggle with visual elements that rely on small fonts, thin borders, or subtle colour differences. Here, multimodal fusion is key: the system can combine direct interaction data with profile information to infer when visual feedback is insufficient, then automatically apply high-contrast themes, larger font sizes, or additional visual cues such as highlighted borders. Because the framework operates in real time, these adaptations can be triggered dynamically when needed, rather than requiring a static accessibility mode to be enabled manually.

Hands-Free Users Some scenarios, whether due to physical impairment, temporary constraints (e.g., carrying items), or task context (e.g., sterile environments) require interfaces that can be operated without direct touch. For these users, the framework supports voice, gesture, and potentially gaze-based interactions. Multimodal fusion allows these input channels to work in combination, reducing ambiguity and improving accuracy. For example, a voice command like “turn on the lamp” can be paired with a pointing gesture to confirm the target device, speeding up the interaction and avoiding false positives.

Why These Groups? These three categories were selected because they cover a spectrum of accessibility challenges that benefit significantly from adaptive, multimodal design: precision (motor), perception (visual), and modality flexibility (hands-free). Each presents distinct technical requirements for sensing,

reasoning, and adapting the interface, making them strong drivers for evaluating the framework’s capabilities. By focusing on these groups, the system demonstrates how a single architectural approach can address varied accessibility needs while remaining extensible to other user categories in future work.

3.2.2 Accessibility Grounding and Adaptation Rationale (WCAG 2.1/2.2)

SIF’s adaptation catalog is not arbitrary. Each action is chosen to operationalise specific WCAG success criteria for motor, visual, or hands-free use. This anchors the framework in established guidance and keeps it portable across platforms. Table 3.1 lists the core adaptations, their intent, and the criteria they support.

Table 3.1: Adaptations mapped to WCAG criteria and intent.

Adaptation	Rationale / User Need	WCAG Support (2.1/2.2)
Increase button size / hit area	Reduces miss-taps for motor-impaired users; improves target acquisition on touch	2.5.5 <i>Target Size (Enhanced)</i> (2.1); 2.5.8 <i>Target Size (Minimum)</i> (2.2)
Increase contrast (UI + controls)	Improves legibility and component discernibility for low vision	1.4.3 <i>Contrast (Minimum)</i> ; 1.4.11 <i>Non-text Contrast</i>
Increase text size / scale	Supports users with low vision and reading difficulty without loss of content/-function	1.4.4 <i>Resize Text</i> ; 1.4.10 <i>Reflow</i> (indirectly, via responsive layouts)
Highlight focus / outline targets	Makes focus location and actionable elements clear for keyboard/assistive tech	2.4.7 <i>Focus Visible</i> ; 2.4.13 <i>Focus Appearance (Minimum)</i> (2.2)
Switch to voice mode (hands-free)	Provides an alternative input when touch precision is poor or hands are unavailable	2.1.1 <i>Keyboard</i> (alternative input availability); 2.5.1 <i>Pointer Gestures</i> (avoid complex gestures)
Simplify layout / increase spacing	Reduces cognitive and motor load; prevents accidental activation of adjacent controls	Supports 1.3.2 <i>Meaningful Sequence</i> ; contributes to 2.5.8 target spacing (2.2)
Tooltips / contextual hints	Aids discoverability and error recovery without relying solely on vision or precision	3.3.1 <i>Error Identification</i> (indirect); general usability aid aligned with WCAG intent

How this is used in SIF. Profiles and recent history decide *when* to apply an adaptation; WCAG mapping decides *which* adaptation is appropriate and safe by default. For example, repeated miss-taps near a small control will raise hit area (2.5.5/2.5.8) and optionally increase contrast (1.4.11) if the profile flags low vision. Hands-free profiles lower the threshold for `switch_mode: voice` to satisfy alternative input availability (2.1.1).

3.3 Frontend Layer: UI Design and Interaction

The Frontend Layer is the user-facing component of the framework, responsible for rendering an adaptive and personalized interface that adjusts dynamically to user needs. Designed for cross-platform deployment, it captures multimodal interactions such as touch, keyboard, voice, or gestures and applies adaptations returned from the backend in real time.

In this thesis, the frontend is demonstrated through the Adaptive Smart Home Controller, a deliberately chosen example that offers both familiarity and a variety of interaction types. The smart home context provides a set of clear, relatable tasks such as switching on a light or adjusting a thermostat that can be adapted for users with motor, visual, or input-related impairments. This use case also highlights the framework’s potential for deployment in other domains where accessibility is a priority.

3.3.1 Interface Elements

The frontend incorporates a set of core UI elements chosen to balance simplicity with the ability to demonstrate a wide range of adaptations. Together, they form a complete interface that could plausibly be used in real-world scenarios, while remaining portable across platforms. Conceptually, these elements can be grouped into three categories:

1. **Action controls:** for example buttons or sliders, which allow users to manipulate device states or settings.
2. **Information displays:** such as text labels, which convey device status or contextual feedback.
3. **Organisational elements:** such as cards or list views, which group related controls for clarity and ease of navigation.

These categories were selected because they cover the most common interaction and accessibility challenges. Action controls benefit from size and spacing adjustments for motor-impaired users, information displays can be adapted with text resizing or high-contrast themes for visually impaired users, and organisational elements help reduce cognitive load by presenting related controls together.

3.3.2 Adaptation Levels

For this thesis and its implementation, adaptations in the frontend can be understood at three conceptual levels:

1. **UI-Level Adaptations:** which modify the appearance of visible elements to improve clarity or ease of interaction.
2. **Geometry-Level Adaptations:** which adjust layout and spacing to reduce input errors or simplify navigation.
3. **Input-Level Adaptations:** which alter the way users interact with the interface, for example by switching to a voice-driven mode.

While the precise mechanics of how these adaptations are implemented, triggered and applied are described in Chapter 5, the design principle remains the same: each change should be both functional and clearly communicated to the user, reinforcing trust and improving interaction efficiency. These high-level principles guide the implementation of the SIF framework, ensuring that adaptations are user-centered and context-aware.

3.4 Input Adapter Layer: Multimodal Input Processing

The Input Adapter Layer acts as the middleware between the user-facing frontend and the reasoning backend, it takes care of bidirectional communication between both layers. Its main role is to take raw interaction data from any modality such as touch, keyboard, voice, or gestures and transform it into a standardised format that the backend can process consistently and vice versa for the frontend. By separating input capture from input interpretation, this layer allows the rest of the system to operate independently of how the input was generated, making it easier to add new modalities in the future (see Figure 3.2 for high-level flow diagram).

Central to this layer is the **JSON Event Contract**, architecturally this is exposed through the **Adaptive UIAdapter** class, which defines the public interface for sending interaction data to the backend. The primary entry point is the `sendEvent(Event eventData)` method. This method accepts an **Event** object, the framework's internal representation of a user interaction, and is responsible for enriching it with contextual information before forwarding it to the backend.

The **Event** class acts as the architectural contract for interaction data. While its exact schema is defined in the implementation (Chapter 5), at the design level it contains:

- **Event type:** the category of interaction, such as a tap, miss-tap, slider miss, or voice command.
- **User ID:** a unique identifier for the user interacting with the system.
- **Timestamp:** the time when the event occurred.
- **Source modality:** the origin of the event (touch, keyboard, voice, gesture, etc.).
- **Target element:** the UI element or control associated with the event.
- **Coordinates:** spatial information where applicable, for example the location of a tap.
- **Confidence score:** a value indicating the certainty of the detected intent.

- **Metadata:** optional context such as UI element type (e.g., button, slider) and more.

By enforcing this structure, the adapter guarantees that all events passed to the backend follow the same rules, whether they come from Flutter, SwiftUI, or a future Unity-based frontend.

The adapter also handles profile verification and creation before sending events. This step ensures that the backend always has a corresponding user profile for contextual reasoning. If the profile does not exist, the adapter triggers its creation using default accessibility settings or a new user profile. These profiles are also bound by a JSON contract to ensure consistency, this design is described in Chapter 4. Furthermore, for low-latency adaptation feedback, events are transmitted via a persistent WebSocket connection, while non-real-time operations such as profile updates, HTTP endpoints are used instead.

Lastly, the adapter manages adaptations returned by the SIF backend layer by listening to the WebSocket channel. When the backend sends adaptation instructions, which is represented in its own JSON contract (in more detail described later), the adapter maps it to an internal representation based off of the JSON contract that can be easily processed by the frontend. This Adapter pattern design is deliberately chosen and platform-agnostic. Although the current implementation is in Dart for Flutter, the architectural pattern can be replicated in any language or platform by implementing the same `sendEvent()` contract and adhering to the same event object structure. This separation means that adding a new modality, such as eye tracking or brain-computer interfaces, only requires implementing the modality's input capture and mapping it to the existing `Event` format, no changes are needed to the backend or reasoning logic.

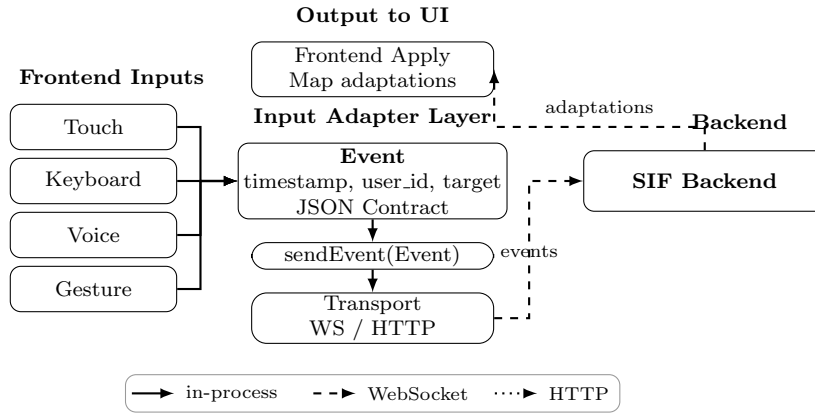


Figure 3.2: High-level flow of the Input Adapter Layer

3.5 SIF Backend Layer: Smart Intent Fusion (SIF)

The SIF backend layer is the reasoning core of the framework, responsible for turning raw user interactions and contextual information into targeted UI adaptations. Operating behind the input adapter, it receives events in the standard JSON contract format and turns them into an internal structured representation (`Event` object), combines them with the user's profile and recent interaction history. This forms the basis for the LLM prompts, next it determines the most appropriate changes to apply to the interface using LLM reasoning.

From an architectural perspective, this layer has two main responsibilities:

- **Processing and interpretation:** validating incoming events, interpreting their intent, and prioritising them according to context.
- **Adaptation generation:** producing a set of structured adaptation actions in a JSON contract that the frontend can apply directly.

The backend is designed to support multiple reasoning strategies. In its current form, it combines deterministic rules with a multi-agent LLM process called Smart Intent Fusion. Rules handle straightforward accessibility needs; for example, increasing button size after a miss-tap, while the LLM process enables more context-aware adaptations that consider multiple factors simultaneously.

Although the underlying logic is covered in detail in Chapter 4, the architectural position of this layer is central in this framework: it serves as the decision-making hub, fed by standardised events from the input

adapter, and returning validated adaptations to the frontend in near real time. This separation allows the reasoning configuration to be changed easily; for example, by swapping models, refining prompts, or integrating new agents without changing the structure of the rest of the framework.

3.6 User Profiles and Context Modeling

User profiles form the backbone of the framework’s personalisation capability. They store accessibility preferences, interaction patterns, and contextual data that allow the system to adapt the interface to an individual’s needs over time. The profile is not a static record, it evolves as the user interacts with the system, incorporating both explicit configuration and implicit observations from their behaviour.

Architecturally, the profile contains three types of information:

- **Static attributes:** such as preferred font size, contrast settings, or dominant input modality, which may be set during initial onboarding.
- **Learned preferences:** derived from patterns in the user’s interactions; for example, frequent miss-taps on small controls may trigger a persistent increase in their size.
- **Contextual data:** including recent interaction history and device environment details, which help the backend reason about the most appropriate adaptations in a given moment.

When a new event is processed, the backend combines it with the relevant profile and context data before passing it to the reasoning logic. This ensures that adaptations are not just reactive to the most recent input, but also informed by longer-term patterns and situational factors.

The architecture treats profiles as a shared resource between all layers:

- The **input adapter** ensures that the correct user profile is referenced with every event, as well as creating or updating the profile when necessary.
- The **backend** updates profiles automatically as adaptations are applied or feedback is received.
- The **frontend** can query profile data to adjust default UI settings before any adaptations are applied.

By integrating profile and context information into every stage of the adaptation pipeline, the framework moves beyond static changes and supports continuous, data-driven personalisation. The specific data schema, storage mechanisms, and update strategies are described in detail in Chapters 4 and 5.

3.7 Dynamic Adaptation Mechanisms

The dynamic adaptation mechanisms are the part of the framework responsible for translating reasoning outputs from the backend into real-time, personalized changes for the user interface. Operating within the SIF backend layer, they take standardised events from the input adapter, combine them with profile and context data, and decide on the most suitable adaptation. These decisions are returned as structured actions to the input adapter, which maps the JSON contract to an internal representation that is then applied immediately by the frontend using callbacks.

At the architectural level, the adaptation process follows a continuous feedback loop. Every interaction is captured, processed, and logged, with the resulting adaptations influencing how the interface behaves in the future. This allows the system to progressively refine its responses, ensuring that changes are not only reactive but also shaped by longer-term usage patterns.

3.7.1 Adaptation Pipeline

The adaptation pipeline follows a structured lifecycle from input capture to UI update, and can be viewed as four main stages (Figure 3.3):

1. **Event reception and context gathering:** Inputs from any modality are captured, standardised, and sent to the backend. The backend retrieves the associated user profile and recent history to build a complete context for reasoning.

2. **Intent interpretation:** The backend analyses the combined event and context, using both deterministic rules and LLM-based reasoning to infer the most likely user intent.
3. **Adaptation Generation:** A set of structured actions is produced, describing changes to be applied to the interface. These may involve modifying the visual presentation, adjusting interaction geometry, or altering input modes. SIF employs:
 - **Rule-Based Logic:** Deterministic rules provide fast, reliable adaptations for common scenarios (mostly used for a reliable backup).
 - **LLM Reasoning:** The AI model infers complex intents and generates creative adaptations for the user.
 - **Heatmap Analysis:** Simulated via history counts, prioritizes problems with frequently interacted elements (e.g., repositioning elements after multiple taps).
4. **Application and logging:** The frontend applies the adaptations immediately, while the backend logs both the event and its outcome for future learning.

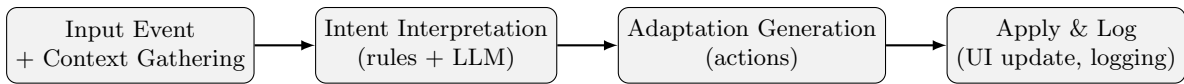


Figure 3.3: Adaptation pipeline flow.

3.7.2 Supported Adaptation Actions

The framework supports a defined set of adaptation types that cover common accessibility needs. These include scaling or spacing adjustments for improved touch accuracy, font and contrast changes for better readability, and modality switches for hands-free interaction. The action set is deliberately kept broad enough to handle varied contexts, yet constrained enough to ensure reliable rendering across platforms (more details in Chapter 5).

3.7.3 Continuous Learning and Feedback Loop

A central design principle is that adaptations are not static changes. The system maintains a history of recent interactions for each user, enabling it to identify recurring patterns such as frequent miss-taps or repeated modality switches. When such patterns are detected, the backend can suggest persistent adjustments; for example, permanently enlarging a frequently used UI element, reducing the need for smaller repeated actions.

3.7.4 Design Considerations

- **Accessibility Focus:** All supported actions are chosen to address motor, visual, and input-related impairments, ensuring that adaptations enhance rather than complicate interaction.
- **Real-Time Performance:** Low-latency communication ensures that adaptations are applied quickly enough to feel seamless.
- **Reliability:** Rule-based logic ensures that adaptations continue to function even if LLM-based reasoning is temporarily unavailable.
- **Extensibility:** The action set and event format are flexible enough to incorporate new adaptation types and modalities in future deployments.

By combining immediate, context-aware changes with a persistent feedback loop, the dynamic adaptation mechanisms give the framework its ability to evolve alongside the user's needs, making it more effective over time.

3.8 Chapter Summary

This chapter has presented the system design and architecture of the multimodal AI-driven framework for dynamic UI adaptation, focusing on its ability to deliver personalized, accessibility-oriented solutions for

motor-impaired, visually impaired, and hands-free users. The framework is built around a modular three-layer architecture: the Frontend Layer, Input Adapter Layer, and SIF Backend Layer that enables the seamless integration of multiple input modalities, including touch, voice, and gestures, and the delivery of real-time UI adaptations.

The architectural design emphasises modularity, scalability, generalisability, and accessibility, creating a solid foundation for extension into other domains such as healthcare and gaming. Communication between layers is handled through WebSocket for low-latency updates and HTTP for reliable profile and debugging operations, while MongoDB and a standardised JSON contract ensure scalability and flexibility in storing and processing interaction data. Together, these elements provide a framework capable of addressing real-world accessibility needs today.

Chapter 4

Smart Intent Fusion (SIF)

4.1 Introduction to Smart Intent Fusion

Smart Intent Fusion (SIF) is the central intelligence layer of the proposed multimodal AI-driven UI adaptation framework, responsible for transforming raw, heterogeneous user input signals into concrete, context-aware interface adaptations. Where the Frontend Layer renders the UI and the Input Adapter Layer standardizes events, SIF performs the reasoning step. Architecturally, SIF occupies the central position in the adaptation loop, receiving standardised events from the Input Adapter Layer, combining them with user profile and context data, and returning validated adaptations back to the frontend for immediate application. This placement ensures that every adaptation decision is both context-aware and modality-agnostic, allowing the framework to maintain consistent behavior regardless of how the interaction was initiated. It fuses current interaction data with user profiles, accessibility requirements, and recent interaction history to infer the user’s underlying intent and translate this into actionable UI changes.

The motivation for SIF stems from a simple but critical challenge in accessibility-focused HCI: users rarely interact with a system through a single, perfectly clean input channel. Instead, interactions are often multimodal, noisy, and incomplete. A motor-impaired user may miss-tap a button but also issue a supporting voice command. A visually impaired user may attempt to activate a control by gesture but with low confidence, relying on high-contrast cues to complete the action. Traditional rule-based adaptive systems tend to process these signals independently, missing the opportunity to combine them into a unified, more reliable understanding of the user’s goal.

SIF addresses this gap through a hybrid reasoning approach:

- Rule-based logic handles deterministic adaptations for example, “if miss_tap on target → increase size by 1.5×” ensuring baseline accessibility support and fast response times even without AI availability.
- LLM-driven reasoning can process richer multimodal context and can propose creative or proactive adaptations that go beyond fixed rules such as switching to voice mode after repeated input struggles, or combining voice + gesture input to trigger a button instantly.
- Multi-Agent SIF (MA-SIF) extends this further by distributing reasoning across specialised LLM agents (UI, Geometry, Input or more) and validating results through a dedicated Validator Agent to reduce hallucinations and conflicting actions as well ensuring the necessary validation.

The SIF layer sits behind well-defined APIs:

- WebSocket endpoint `/ws/adapt` for low-latency, real-time adaptation suggestions.
- HTTP endpoints like `/profile` for profile management, or `/full_history` for developer tooling and debugging.

These endpoints use a strict JSON contract, making it easy for any frontend platform to connect and work with the backend. This ensures that the system remains flexible and can handle different input modalities.

A defining feature of SIF is its integration with persistent user profiles stored in MongoDB. These profiles encapsulate three layers of information: static attributes such as preferred font size or contrast mode, learned preferences derived from recurring interaction patterns, and contextual data including recent history and environmental details. By merging these dimensions with each incoming event, SIF can make decisions that are not only responsive to the user’s current action but also aligned with their long-term accessibility needs. By fusing this profile data with incoming events, SIF maintains a continuous personalization loop progressively adapting the UI to match the user’s abilities and context over time.

For example:

- A profile with `motor_impaired: true` will bias adaptations towards larger controls and simplified layouts.
- Repeated history of slider misses may lead to permanent slider thumb enlargement.
- A hands-free preference can automatically promote voice/gesture-driven navigation in relevant contexts.

In the context of this thesis, SIF is not only an internal backend feature, it is the core research contribution. The remainder of this chapter expands on the theoretical underpinnings of SIF, its integration with user profiles and multimodal fusion, the prompt engineering strategies for guiding LLM behavior, the architecture of its multi-agent extension, and the performance metrics used to evaluate its effectiveness.

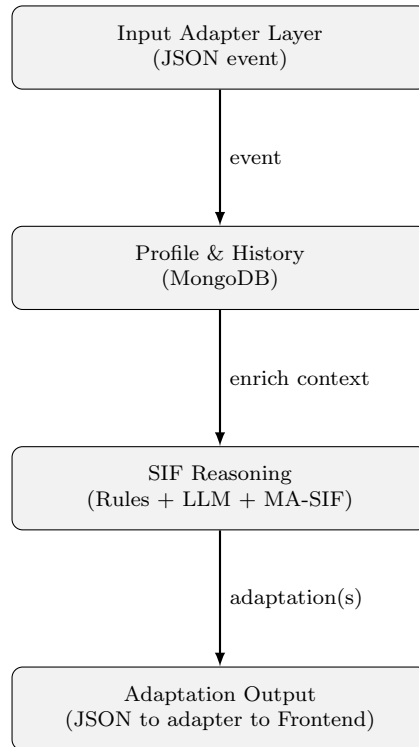


Figure 4.1: Abstract event flow in Smart Intent Fusion: events from the Input Adapter are enriched with profile/context, processed by rule-based and LLM/MA-SIF reasoning, and returned as adaptation outputs to the frontend.

4.2 Theoretical Foundations of Smart Intent Fusion

The idea behind Smart Intent Fusion is simple:

users don’t interact with a UI in one clean and perfect way. In real life, inputs are messy, mixed, and often incomplete. A person might tap the wrong spot on a button, then say a voice command to make sure it worked. Someone using gestures might point at something but not be perfectly accurate, so they rely on extra visual cues to finish the task. A lot of current adaptive systems still process these signals

separately, and that means they miss the chance to combine them into a clearer picture of what the user actually wanted.

Smart Intent Fusion tries to fix that by taking all the inputs together: touch, keyboard, voice, gestures and mixing them with what the system already knows about the user (their profile, their history of interactions, and their accessibility needs). This way, it can guess the real intent and adapt the UI in the most helpful way.

4.2.1 Multimodal Fusion

In HCI, “multimodal fusion” just means combining different types of input to get a more reliable or richer understanding of the user’s action.

This can be done in a few ways:

- **Early fusion:** merge the raw signals before interpreting them (e.g., combining the coordinates of a tap with the audio of a voice command immediately).
→ Useful when two inputs happen at the exact same moment.
- **Late fusion:** process each input type separately first, then merge the interpreted results (e.g., “voice command = unlock” + “miss-tap near unlock button” → final decision = trigger Unlock button).
→ This is where most SIF reasoning happens.
- **Hybrid fusion:** a mix of both, sharing some data early but also combining results later.
→ Used when one modality’s data can make another modality’s interpretation more accurate.

SIF uses the hybrid approach. The Input Adapter Layer already standardizes each input into a clean JSON format, but the fusion step happens in the backend where the current event, past history, and profile are all processed and fused together. This makes it possible to combine patterns like “miss-tap + voice command” into one clear adaptation. In the overall architecture described in Chapter 3, this fusion step is the bridge between raw input handling and adaptation generation. By sitting in the backend, SIF can remain entirely modality-agnostic while still benefiting from the structured event format defined by the Input Adapter Layer. This ensures that even when new input methods are introduced in the future, the fusion logic can remain unchanged, relying on the same JSON schema and profile-context pipeline.

4.2.2 Intent Inference

The main goal of SIF is not just to log inputs, but to figure out *why* the user did them. This is called intent inference. In other words, we want the system to answer the question: “What was the user trying to do?” If the system knows the intent, it can choose the best adaptation. The inference process is always profile-aware. Each decision is informed by a combination of static preferences (such as font size or contrast mode), learned patterns from past usage (such as repeated difficulty with sliders), and short-term contextual data (such as the most recent interactions). This means that two users producing the same input sequence could receive different adaptations if their profiles and histories differ, keeping the interaction completely personalised.

For example:

- If the intent is to press “Unlock” but the user misses, the adaptation could be to enlarge the button and trigger it right away.
- If the intent is to adjust a thermostat but the user struggles with the slider, the adaptation could be to switch to voice control and increase the slider size.

This is where LLMs can help. They are good at reasoning about context, combining clues, and filling in gaps when inputs are unclear. The downside is that they can be slow, be costly to run, and sometimes “hallucinate” an answer that doesn’t make sense even if it’s not defined in the prompt. That’s why SIF uses a **hybrid** approach:

- **Rules** handle clear, simple cases with instant feedback.
- **LLM reasoning** handles more complex or ambiguous cases.

4.2.3 Why Hybrid Works Best

A purely rule-based system is predictable but rigid. It can only respond to situations that were thought of in advance. A purely LLM-based system is flexible but not always reliable, especially when it needs to work in (close to) real-time for accessibility. By combining both:

- The rules guarantee that basic accessibility changes (like increasing size after a miss-tap) always work. For example:

```
IF event_type == "miss_tap" AND profile.motor_impaired == true
THEN action = increase_button_size(target, 1.5)
```

- The LLM adds creativity and can adapt to situations the rules didn't cover, like combining unusual input patterns, proposing a mode switch for hands-free use or even something entirely new.

4.2.4 Connection to Accessibility

The whole point of this is to improve accessibility for different kinds of users:

- **Motor-impaired:** combine multiple input signals to avoid repeated failed attempts.
- **Visually impaired:** recognize when visual feedback is not enough and trigger higher-contrast or bigger fonts automatically.
- **Hands-free:** allow combinations like voice + gesture to instantly activate actions.

In short, the theoretical base for SIF comes from multimodal fusion, intent inference, and hybrid reasoning. These ideas are not new in HCI, but this framework applies them with a strong focus on accessibility and personalization, and makes them work in real-time with cross-platform UI code. For example, a motor-impaired user who misses a lock button twice and then issues a voice command might trigger a combined adaptation: the UI immediately enlarges the button for future taps, but also switches the interface to voice-first mode for the current session. This ability to layer short-term fixes on top of long-term adjustments is what makes SIF more than just a reactive system, it is a continuously learning adaptation layer.

4.3 User Profile and Context Integration

For Smart Intent Fusion to be truly “smart,” it needs more than just the current event it is processing. If SIF reacted to every tap, voice command, or gesture without knowing who the user is or how they usually interact, it would behave like a generic accessibility script rather than a personalised adaptation system. That's why the user profile and interaction context form the backbone of the reasoning process. They give SIF a sort of memory/personality, and the ability to adapt over time, not just in the moment. When a new event comes in like a tap, miss-tap, voice command, or gesture, SIF doesn't look at it in isolation. It combines that event with:

1. **The user profile** – a stored record in MongoDB with accessibility flags, preferred modalities, and UI settings.
2. **Interaction history** – the last 10 events that show patterns or repeated problems.
3. **Current UI context** – optional metadata about what's on screen, where buttons are placed, and their sizes.

This means SIF can make decisions that are personal and context-aware, not just reactive.

4.3.1 User Profiles

A user profile stores the information that makes one person's interaction style different from another's. This can include accessibility needs (motor-impaired, visually impaired, hands-free preferred), preferred input modalities, and baseline UI settings like font size, contrast mode, and button scale. It can be seen as the **memory** of the system. When combined with a short history of recent interactions, this profile turns SIF from a static decision engine into a continuous learning system.

Without profiles, SIF could still make adaptations, however they would always be reactive and temporary. For example:

- If a user with tremors keeps missing a button, the button might get enlarged for that session, but as soon as they restart the app, it would shrink back.
- If a user prefers voice input, SIF wouldn't know to automatically switch to voice mode when they struggle with touch, instead it would have to come up on its own that this switch is needed for this user.

Profiles ensure these adaptations stick and get better over time.

4.3.2 User Profile Structure

In the backend, each profile is a JSON document in the `profiles` collection of a MongoDB database. It's indexed by `user_id` so the system can look it up instantly whenever a new event arrives. A typical structure looks like this:

```

1 {
2   "user_id": "user_123",
3   "accessibility_needs": { "motor_impaired": true },
4   "input_preferences": { "preferred_modality": "voice" },
5   "ui_preferences": { "font_size": 16 },
6   "interaction_history": [
7     { "action": "increase_button_size", "target": "all", "value": 1.2 },
8     { "action": "increase_font_size", "target": "all", "value": 1.1 }
9   ]
10 }
```

Listing 4.1: Simplified User Profile Example

This design keeps it simple but flexible. Architecturally, each profile contains:

- `accessibility_needs`: flags that tell SIF what kind of adaptations to prioritize.
- `input_preferences`: helps the system decide which modality to switch to when needed.
- `ui_preferences`: baseline visual parameters such as font size and contrast mode.
- `interaction_history`: capped log of recent events to support continuous learning.

MongoDB's indexing means the profile can be retrieved in milliseconds, even with a large user base, and capped histories ensure lookups are fast.

4.3.3 How Profiles Affect Decisions

When an event comes in, the backend follows a clear process:

1. Load the profile from MongoDB using the `user_id`. If it doesn't exist yet, create a new default profile.
2. Combine the event with the last few interactions from the history.
3. Pass the profile, history, and current event into the Smart Intent Fusion reasoning step.

This context completely changes how SIF decides on adaptations.

Some examples of influenced decisions:

Example 1: Motor-impaired user with repeated miss-taps

If the last three events in history are miss-taps on the same “Unlock/Lock” button, and `motor_impaired` is true in their profile, SIF might do two things at once:

```

1 [
2   {"action": "increase_button_size", "target": "button_unlock", "value": 1.5,
    ↪ "reason": "Multiple miss-taps detected, enlarging button for better
    ↪ accessibility"},
3   {"action": "highlight_border", "target": "button_unlock", "reason": "Increase
    ↪ button visibility for the user"}
4 ]

```

Listing 4.2: Possible motor-impaired user adaptations

Without the profile, it might have just enlarged the button once and moved on.

Example 2: Hands-free preferred user A user with `"hands_free_preferred": true` points at a device card (gesture) and says “turn on the lights.” SIF reasoning can fuse these into:

```

1 {
2   "action": "trigger_button",
3   "target": "button_light",
4   "reason": "Gesture pointing + voice 'Turn on the lights' detected for
    ↪ hands-free user",
5   "intent": "Activate Lights"
6 }

```

Listing 4.3: Hands-free user intent fusion

Because of the profile, SIF is confident enough to trigger the button immediately without asking for physical confirmation.

4.3.4 Continuous Learning from History

Profiles are not static. Every interaction is logged in the `interaction_history` and can influence future decisions. The interaction context can be seen as the system’s **short-term awareness**. This makes SIF a learning system:

- If a user keeps manually enabling high-contrast mode, the system can update `contrast_mode` in their profile so it’s always on by default.
- If increasing a button size significantly reduces miss-taps, that size can become the new permanent baseline in `ui_preferences`.
- If switching to voice mode solves repeated touch struggles, the profile can be updated to favor voice by default.

This feedback loop means the user doesn’t need to “train” the system manually, it adapts naturally as they use it. If SIF only looked at the current event, it would miss important patterns and make short-sighted decisions. With profile + history + event combined:

- Adaptations can be proactive instead of reactive.
- The UI can stay consistent between sessions.
- The system can learn what really helps the user over time.

From an accessibility perspective, this is the difference between a generic interface that occasionally helps and a personalised tool that feels like it understands the user.

4.3.5 Role in Accessibility

User profiles are the backbone of accessibility-focused adaptations in this framework. They act as a persistent memory of the user’s abilities, preferences, and interaction challenges, enabling the system

to make targeted, proactive adjustments rather than relying solely on short-term reactive changes. By storing explicit accessibility flags alongside learned behavioral patterns, profiles allow SIF to tailor the interface to an individual user in ways that are both short-term and long-term.

From an accessibility perspective, profiles influence SIF’s reasoning in three key user categories:

- **Motor-Impaired Users:** Profile flags such as `motor_impaired: true` prioritize adaptations that reduce fine motor precision requirements. This can include:
 - Enlarging touch targets (buttons, sliders) and increasing spacing to prevent accidental taps.
 - Offering alternative modalities such as voice commands or keyboard navigation to bypass touch interaction altogether.
 - Retaining enlarged target sizes across sessions once repeated miss-taps are detected.
- **Visually Impaired Users:** When `visual_impaired: true` is set, adaptations aim to maximise visual clarity. Examples include:
 - Switching to high-contrast themes and bold color schemes to align with WCAG 2.1 contrast requirements.
 - Increasing font sizes and icon scales to meet text accessibility guidelines.
 - Highlighting the active element with a strong focus border or magnified overlay to improve navigation feedback.
- **Hands-Free Users:** Profiles with `hands_free_preferred: true` bias SIF towards non-touch input modes, reducing physical interaction demands. Adaptations may involve:
 - Automatically switching to voice or gesture navigation when interaction struggles are detected.
 - Providing clear, speech-friendly UI labels and tooltips to improve command recognition accuracy.
 - Simplifying layouts by reducing the number of visible controls at once, making it easier to select elements through voice or gesture.

By embedding these accessibility considerations directly into the profile structure, SIF can reason in a way that is both context-aware and user-specific. This allows the system to:

1. Anticipate needs before errors occur (e.g., pre-emptively enlarging critical controls for a motor-impaired user on a small screen).
2. Ensure adaptations persist across sessions, avoiding the frustration of having to reconfigure accessibility settings each time.
3. Combine profile knowledge with real-time interaction patterns, enabling nuanced decisions such as *“keep high contrast on by default, but also enable voice mode when the user is multitasking”*.

4.4 Modeling Multimodal Input Fusion

Smart Intent Fusion doesn’t just take one input, it collects **multiple inputs from different modalities** like touch, voice, gestures, and more. Then fuses them into a single, well-reasoned adaptation. This process is called *multimodal input fusion*. Without fusion, the system would treat each event separately. A miss-tap would trigger one adaptation, and a voice command would trigger another, without realising both were aimed at the same action. With fusion, those two inputs can be combined into one confident and more helpful response which saves time and reduces frustrations. Furthermore, the user profile and interaction history also directly affect how this fusion works. A hands-free preferred user will have a lower threshold for fusing gesture + voice, but a visually impaired user’s profile might cause SIF to always add a “highlight” adaptation when triggering elements via voice, even if not strictly necessary. If history shows repeated failures for a certain modality, its weight in fusion decisions can be reduced. For users with impairments, every extra action is extra effort. Because of how the fusion is designed, it cuts down on unnecessary steps so the UI adapts faster and smarter to the user’s needs.

4.4.1 Event Standardisation

Before any fusion can happen, the raw input needs to be standardised. Every frontend in the framework Flutter, SwiftUI or a future VR client converts its local input data into the same JSON contract. This is handled by the Input Adapter Layer as described earlier.

A standardised event looks like this:

```

1 {
2   "event_type": "miss_tap",
3   "source": "touch",
4   "timestamp": "2025-08-04T14:41:00Z",
5   "user_id": "user_123",
6   "target_element": "button_unlock",
7   "coordinates": {"x": 210, "y": 640},
8   "confidence": 0.8,
9   "metadata": {"UI_element": "button"}
10 }
```

Listing 4.4: Standardised Event Example

This contract ensures:

- **Cross-platform compatibility:** all clients speak the same “language” to the backend.
- **Dynamic field names:** fields like `event_type` and `target_element` can be extended and modified without breaking existing logic.
- **Extensibility:** new modalities can be added (like gaze tracking) without breaking existing clients.

By the time the event reaches SIF, it doesn’t matter whether it came from a phone, desktop, or VR headset, it always adheres to the same format for the backend. As said earlier SIF uses a hybrid approach for multimodal input fusion. The Input Adapter handles basic pre-processing (similar to early fusion), but the actual reasoning, deciding what adaptations to apply happens in the backend using late fusion.

Example: Touch + voice

Let’s say a motor-impaired user taps just to the right of the “Lock/Unlock” button and, within a second, says “Unlock.” Individually, the miss-tap could trigger an enlargement of the button and the voice command could trigger the unlocking action. However, with multimodal fusion, the system can understand that both inputs are related and prioritize the unlocking action while also enlarging the button for better accessibility.

With SIF fusion, the system sees both in context:

```

1 [
2   {"action": "increase_button_size", "target": "button_unlock", "value": 1.5,
3     ↳ "reason": "Miss-tap detected near Unlock button"},
4   {"action": "trigger_button", "target": "button_unlock", "reason": "Voice
5     ↳ command 'unlock' detected in combination with miss-tap"}
6 ]
```

Listing 4.5: Example of Multimodal Fusion

Now the button is both enlarged for future use and triggered immediately, reducing the number of actions the user needs to take, simplifying the interaction which is useful for motor-impaired users.

4.4.2 Timing and Confidence

SIF’s LLM reasoning can also consider **when** and **how confidently** an input happened based on two fields from the standardised event:

- **Timing:** events close together in time (e.g., within 1–2 seconds) are more likely to be related.
- **Confidence:** each modality can provide a confidence score (e.g., gesture detection might be 0.7 certainty). Lower confidence might require a second modality before acting.

For example, a low-confidence gesture to point at a button might do nothing alone, but if followed by a high-confidence voice command naming that button, SIF can treat them as a combined intent.

4.4.3 LLM Reasoning in Fusion Decisions

While the fusion process benefits from deterministic rules for speed and reliability, one of the key innovations of SIF is its ability to leverage LLM reasoning to interpret and combine multimodal inputs in a context-aware way.

Once standardised events reach the backend, they are not processed in isolation. Each event is enriched with:

- **User profile data:** accessibility flags, preferred modalities, and baseline UI settings.
- **Interaction history:** the most recent events, revealing repeated struggles or patterns.
- **Contextual metadata:** details about the UI state (e.g., which elements are visible), UI element type, environment and more.

This enriched dataset is then included in the LLM prompt, allowing the model to reason comprehensively about the user’s intent across modalities. For example, instead of treating:

- a low-confidence “point” gesture at a thermostat slider, and
- a voice command “set temperature to 22”

as separate actions, the LLM can infer that they describe the same goal and produce a single adaptation:

```

1 {
2   "action": "adjust_slider",
3   "target": "slider_thermostat",
4   "value": 22,
5   "reason": "Gesture and voice command combined to adjust temperature"
6 }
```

By bringing together different types of input modalities, user’s preferences, and their recent actions, SIF can suggest changes that are not only accurate but also anticipate what will help the user most. This means the system can adapt quickly and make the interface easier to use based on what the user’s most probable intent was.

4.5 Rule-Based Logic and LLM-Driven Adaptation

Smart Intent Fusion uses two very different ways to decide what adaptation to apply:

- **Rule-based logic** handles clear, deterministic cases where the system can apply a known adaptation based on the event type and user profile. It is instant, and predictable.
- **LLM-driven reasoning** uses the Gemini API to process complex, multimodal contexts and propose creative adaptations that go beyond simple rules.

Both have strengths and weaknesses, which is why the framework combines them instead of choosing one.

4.5.1 Rule-Based Logic

Rule-based logic works by matching incoming events to predefined conditions and applying a fixed response.

For example:


```

1   if event.event_type == "miss_tap":
2       return {"action": "increase_button_size", "target": event.target_element, "value": 1.5}

```

Advantages:

- **Simplicity:** Easy to implement and understand.
- **Speed:** Instantaneous responses to known events.
- **Predictability:** Consistent behavior for similar inputs and no risk of unintended consequences like hallucinations.
- **Baseline guarantee:** Acts as a safety net so that critical accessibility features still work if the LLM is unavailable, slow to respond, or returns unusable output.

Limitations:

- Can only handle cases explicitly programmed in advance.
- No ability to combine signals in creative ways.
- Doesn't learn new patterns unless a developer updates the rules.

This means that while rule-based logic is fast and reliable, it can also be rigid and unable to adapt to new situations without human intervention. In SIF, the rules are deliberately kept lightweight, mostly as a mock or backup layer for LLM-driven reasoning.

4.5.2 LLM-Driven Reasoning

The LLM can reason about the event, user profile, and history together to try and infer the user's intent more deeply. It can make connections that rules would miss, such as:

- Combining a miss-tap with a voice command into a single action.
- Switching to a different modality when it detects repeated failure in the current one.
- Proposing multiple coordinated adaptations for one intent.

Advantages:

- **Flexibility:** Can adapt to new situations without explicit programming by the developer.
- **Context-aware:** Takes profile and history into account naturally.
- **Learning:** Can improve over time by learning from user interactions and feedback.

Limitations:

- **Less predictable:** May generate unexpected or irrelevant responses.
- **Slower:** Network call + reasoning time.
- **Needs validation:** Output must be checked more thoroughly before applying.

4.5.3 Hybrid Approach in SIF

In practice, SIF doesn't fully choose between the two, it blends them:

1. **Rules first:** If the event matches a clear, high-confidence rule, apply it immediately, this is mostly done in the frontend by the user profile.
2. **LLM second:** Use the model for complex or ambiguous cases, or to suggest extra adaptations beyond the rules.
3. **Timeout Fallback:** If the LLM times out or fails, return rule-based or other LLM output only.

This ensures that basic accessibility features always work, while still allowing for creative, context-aware adaptations when needed. In other words, the rules form the “floor” of the system (the minimum guaranteed level of accessibility) while the LLM can raise the ceiling by adapting to more complex,

ambiguous, or novel situations. Users are never left waiting for AI responses that might never come, and the system remains responsive even in worst-case scenarios.

4.5.4 Heatmap Analysis

To further refine adaptation decisions, the backend can analyze interaction heatmaps derived from user event logs. By aggregating tap coordinates and gesture paths, the system identifies problematic UI regions (e.g., frequently missed buttons) and adapts layouts or element sizes accordingly. This data-driven approach supports continuous improvement and personalization, especially for users with evolving accessibility needs. By using tap frequency, the backend can suggest adaptations like repositioning elements or enlarging hit areas, enhancing usability.

4.6 Multi-Agent Smart Intent Fusion (MA-SIF)

While a single LLM can process events and suggest adaptations, it often tries to “do everything” at once. That makes it harder to constrain, more prone to hallucinations, less predictable and in some cases slower (depending on workload). Multi-Agent Smart Intent Fusion (MA-SIF) solves this by splitting the reasoning into specialised agents, each focused on one domain of UI adaptation, and then combining their outputs through a Validator Agent. This design brings the benefits of modularity, parallelism, and role-specific constraints, all of which improve reliability and make the system easier to maintain.

4.6.1 Why Multiple Agents?

SIF began with a single-agent LLM that handled UI changes, geometry edits, and input switching in one pass. In practice, this made the prompt broad and brittle, and JSON outputs were less consistent. Measured against the same event suite, the single-agent baseline achieved lower schema validity (cf. Section 4.8) and was harder to constrain. MA-SIF addresses this by dividing work across focused agents (UI, Geometry, Input) and routing all suggestions through a Validator Agent. The result is tighter control over allowed actions, fewer duplicates, and higher parse reliability, at the cost of a small latency increase. Furthermore, when one model is asked to handle UI changes, geometry adjustments, input mode switching, and validation all at once, several problems appear:

- The prompt becomes long and vague.
- Allowed actions become harder to enforce and turn into hallucinations.
- Reasoning gets scattered between unrelated concerns.

By giving each agent a clear role, their prompts can be short, specific, and easy to maintain. For example, the UI agent only ever sees actions like `increase_font_size` or `increase_contrast`, while the Geometry agent only deals with spatial changes like `increase_button_size` or `adjust_spacing`. This separation means that each agent can focus on its specific task without being distracted by unrelated concerns. Parameters can be tuned independently for each agent, allowing for more precise control over their behavior (e.g., lower temperature for Geometry, slightly higher for Input). Finally, debugging becomes easier as each agent’s logic is contained and easier to follow.

4.6.2 Agent Roles

At a high level, MA-SIF in this thesis’s current configuration, consists of four specialised LLM agents:

1. a **UI Agent** for visual and interactive adaptations,
2. a **Geometry Agent** for spatial changes and layout simplification,
3. an **Input Agent** for modality switching and interaction simplification,
4. a **Validator Agent** for conflict resolution and accessibility compliance.

Each has a narrow focus, defined prompts, and a limited set of allowed actions, making their reasoning predictable and easier to validate. They have clearly defined scopes of responsibility, and their roles are deliberately narrow to keep prompts concise, outputs predictable, and debugging straightforward. **The UI Suggestion Agent** is concerned purely with visual accessibility changes such as increasing font size, toggling high-contrast mode, or displaying contextual tooltips. For example, when a visually impaired

user interacts with a dense text block, this agent might output an adaptation like increasing the global font scale by $1.2\times$.

The **Geometry Suggestion Agent** deals with spatial layout changes and the physical dimensions of interactive elements. It might recommend increasing the size of a button, expanding the hit area of a slider, or adding extra spacing between cards in a list. For example, a motor-impaired user who repeatedly misses a button may trigger an output to enlarge that specific button by $1.5\times$ while keeping the rest of the interface unchanged.

The **Input Suggestion Agent** focuses on modality switching and simplifying interaction pathways. It is capable of suggesting transitions between touch, keyboard, voice, or gesture modes, as well as proposing layout simplifications to reduce cognitive load. For instance, if a user struggles with touch but succeeds with voice commands, this agent can suggest switching to voice-first navigation, potentially combined with a reduced interface complexity.

Finally, the **Validator Agent** operates after all others have completed their reasoning. This is the most computationally demanding role, as it must examine every proposed adaptation in detail, identify and remove duplicates, verify that all actions are allowed, and ensure that no output falls outside safe parameter ranges. It also resolves potential conflicts, such as two agents targeting the same element with incompatible values. Because of the breadth of this responsibility, the validator uses a larger Gemini model, a higher timeout of 30 seconds, and a dynamic thinking budget.

4.6.3 Adaptation Flow

When an event arrives at the backend, the fusion process begins by loading the user’s profile and recent interaction history from MongoDB. This contextual data is then sent to each non-validator agent, alongside the event itself. The prompt for each agent is tailored to its domain, ensuring that the LLM only receives relevant instructions and the list of actions it is permitted to output.

Once each agent has processed the event, their suggestions are collected. Importantly, the system is tolerant of partial failures, if one or more agents fail to return a result due to a timeout or API error, the remaining agents’ outputs are still retained. These partial results are then passed to the Validator Agent, which merges them into a coherent and conflict-free final adaptation set. Another possibility is fusing LLM adaptations with rule-based suggestions to create a hybrid output when one or more agents fail.

If the validator itself fails, either due to malformed output or a processing error, the system does not discard all results. Instead, it falls back to returning the raw, unvalidated suggestions from the agents since the output may still be more useful than rule-based alternatives. Only if all agents fail to produce output does the framework revert to the rule-based `mock_fusion` fallback. This layered approach ensures that useful adaptations are preserved whenever possible, rather than being lost due to a single point of failure.

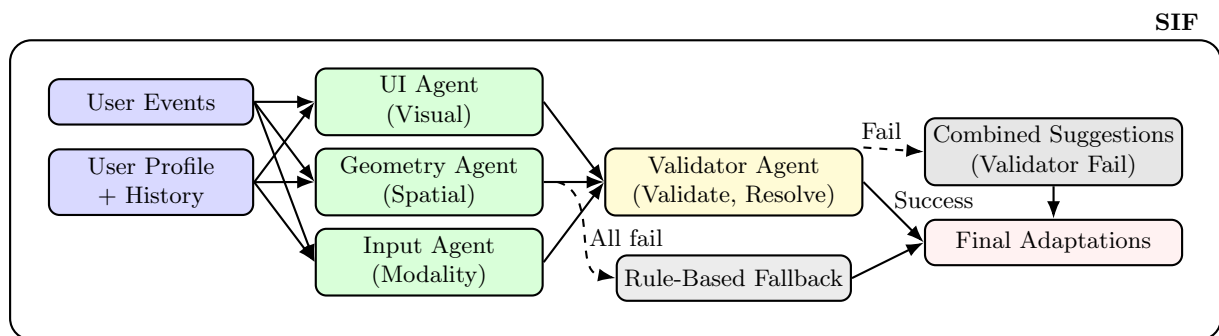


Figure 4.2: Multi-agent SIF adaptation flow: user events and profile/history are processed by specialised agents, validated, and merged into final adaptations. Fallbacks ensure robustness.

4.6.4 Dynamic Configuration

A key advantage of MA-SIF is that it is configured almost entirely through the `sif_config.json`, making it very adaptable without requiring backend code changes. Each agent's behavior can be fine-tuned individually including its allowed actions, model type, temperature, thinking budget, and timeout, simply by editing the configuration file. This flexibility extends to the ability to add new agents or duplicate existing ones with different focuses.

For instance, it is entirely possible to run two Geometry agents in parallel: one optimized for mobile devices with smaller screens and one tuned for desktop layouts. Both would analyse the same event and profile data but apply different heuristics in their prompts. The Validator Agent would then merge their suggestions, resolving overlaps and conflicts automatically. This makes it possible to create specialised agents for emerging modalities, such as gaze tracking in VR, without altering any core fusion logic.

This externalised configuration is particularly valuable in accessibility research, where rapid iteration is needed. For example, during a live user study, the prompt for the UI Agent can be adjusted to emphasise high-contrast themes over font scaling without redeploying the backend. Similarly, experimental agents for new modalities such as gaze tracking or haptic feedback can be added in minutes, allowing researchers or developers to test new adaptation strategies with minimal downtime.

```

1  {
2    "ui_agent": {
3      "model_settings": {
4        "model": "gemini-2.5-flash-lite",
5        "temperature": 0.2,
6        "thinking_budget": 0
7      },
8      "focus": ["visual adaptations"],
9      "allowed_actions": ["increase_size", "increase_contrast",
10     ↪ "reposition_element"],
11     "prompt": "Given the event data, suggest UI adaptations for accessibility.
12     ↪ Consider user profile and interaction history."
13   },
14   "geometry_agent": {
15     "model_settings": {
16       "model": "gemini-2.5-flash-lite",
17       "temperature": 0.2,
18       "thinking_budget": 0
19     },
20     "focus": ["spatial adaptations"],
21     "model": "gemini-2.5-flash-lite",
22     "allowed_actions": ["resize_element", "adjust_spacing", "simplify_layout"],
23     "prompt": "Analyze the UI layout and suggest spatial adaptations to improve
24     ↪ usability for motor-impaired users."
25   },
26   "validator_agent": {
27     "model_settings": {
28       "model": "gemini-2.5-flash",
29       "temperature": 0.3,
30       "thinking_budget": -1,
31       "timeout": 30
32     },
33     "allowed_actions": ["switch_to_voice", "interpret_gesture",
34     ↪ "recover_from_error", "increase_size", "increase_contrast",
35     ↪ "reposition_element", "resize_element", "adjust_spacing", "simplify_layout"],
36     "prompt": "Validate proposed adaptations for conflicts, duplicates and
37     ↪ inconsistencies based on user context, events and interaction history."
38   }
39 }

```

```

32 }
33 }

```

Listing 4.6: Basic example of a 2-agent (+ validator) configuration

Each agent configuration above specifies its model, temperature, and thinking budget, which control the LLM’s behavior and output style. The `focus` field describes the agent’s adaptation domain, while `allowed_actions` restricts the set of suggestions to ensure predictable outputs. The `prompt` guides the agent’s reasoning, and is dynamically extended at runtime with the current event (see Chapter 5), user profile, and interaction history. This enables agents to generate personalized adaptation suggestions based on live user context, rather than relying solely on static instructions.

4.6.5 Temperature and Thinking Budget

Two key parameters control the behavior of each LLM agent in MA-SIF: **temperature** and **thinking budget**.

Temperature determines the randomness and creativity of the model’s output. Lower values (e.g., 0.2) make the agent more deterministic, producing consistent and predictable suggestions—ideal for accessibility-focused UI changes where reliability is critical. Higher values (e.g., 0.7 or 1.0) increase creativity and output diversity, but can introduce unpredictability or hallucinations, which may be undesirable for critical adaptations.

Thinking budget limits the number of reasoning steps or computational effort before the agent responds. A budget of 0 results in minimal reasoning and fast, direct answers. Higher budgets allow the agent to consider more factors, leading to more nuanced or complex suggestions. Negative values (e.g., -1) remove the limit, enabling extensive reasoning, useful for roles like the Validator Agent, which must resolve conflicts and ensure schema compliance.

Effect on Reasoning:

- Lower temperature and budget: Fast, predictable, and simple adaptations.
- Higher temperature and budget: More creative, thorough, and potentially complex adaptations, but with increased unpredictability and latency.

In summary, temperature and thinking budget are critical for tuning agent behavior. Lower values favor speed and reliability, making them suitable for accessibility-focused changes. Higher values enable deeper reasoning and creativity, but may introduce unpredictability, which can be risky for critical UI adaptations. The framework defaults to low temperature and minimal thinking budget for suggestion agents, while the Validator Agent uses a higher budget and slightly increased temperature to ensure robust validation.

4.6.6 Example in Action

Consider a user with both `visual_impaired: true` and `hands_free_preferred: true` in their profile. They attempt to turn on the lights by tapping near the “Turn on” button but miss slightly, then say “Turn on the lights” almost immediately afterwards. The UI Suggestion Agent, informed by the user’s visual impairment, recommends switching to high-contrast mode and displaying a tooltip to make it easier to know the app’s workings next time. The Geometry Suggestion Agent identifies the repeated miss and proposes increasing the button size. The Input Suggestion Agent recognizes the voice command and suggests both switching to voice mode and triggering the button directly.

These suggestions are passed to the Validator Agent, which removes any duplicate `increase_button_size` actions, ensures all parameters are sound, and merges the remaining actions into a single, ordered list. The final adaptation set includes the contrast adjustment, the button enlargement, the tooltip display, the voice mode switch, and the direct triggering of the button. All of these are applied in one update, making the interface immediately more accessible and easier to use in future interactions with the possibility of making them permanent for the user, by updating their profile.

This scenario also highlights MA-SIF’s integration with persistent profiles and recent history: because the profile already records both visual impairment and a hands-free preference, the agents start from a

position of context-awareness rather than guessing from scratch. The resulting adaptations are therefore both reactive to the immediate miss-tap and proactive in aligning with the user’s long-term accessibility needs.

4.6.7 Benefits of the Multi-Agent Approach

The multi-agent architecture provides several practical advantages. Reliability is improved because a failure in one agent does not prevent others from giving valuable suggestions. The modularity of the design makes it straightforward to maintain, as each agent can be modified or tuned independently without risking regressions in unrelated areas. Running agents in parallel also improves responsiveness, particularly when some suggestions can be applied even before all agents have finished processing.

Specialisation further enhances the quality of the output, as each agent’s prompt scope is narrow enough to minimise irrelevant reasoning. The dynamic configuration system makes it possible to scale the number of agents up or down, or to swap in different models, without any changes to the backend logic. This adaptability is especially important in research and prototyping contexts, where requirements may change quickly.

The practical implementation of MA-SIF including the `sif_config.json` schema, example prompts, and backend implementation is detailed in Chapter 5.

4.7 Prompt Engineering for LLMs in SIF

One of the most important parts of Smart Intent Fusion is **how we talk to the LLM**. Unlike a traditional rule-based system, where logic is written explicitly in code, here the behavior of the LLM depends on the instructions it receives, namely the prompt. If the prompt, by which the LLM is queried, is unclear, missing context, or too open-ended, the output will either be wrong, inconsistent, or impossible to parse in code. The LLM could also be more prone to hallucinations, which means it generates answers that sound plausible but are actually incorrect or nonsensical. This means the design of the prompt, the model parameters, and the output format all directly affect how useful and reliable the adaptations are. Even though the prompts used in this framework are relatively simple compared to large fine-tuned systems, they still follow a consistent structure and design philosophy that make them work for this use case. In the architecture described in Chapter 3, each event passed to the SIF Backend Layer arrives already standardised, enriched with metadata, and paired with relevant user profile and history context. The following prompt engineering process simply embeds this structured data into the LLM request, ensuring that reasoning always starts from the same reliable, modality-agnostic representation.

4.7.1 LLM Prompt Design Principles

Prompt engineering in this context is not just about getting a correct answer, it directly impacts accessibility outcomes. A poorly constrained prompt could suggest adaptations that introduce visual clutter, require unnecessary actions, or even reduce usability for the intended audience. By embedding accessibility goals, WCAG criteria, and known user needs into the prompt, the LLM is guided toward changes that genuinely improve the interface rather than merely altering it. The main objectives that flow from this understanding when writing the SIF prompts were:

1. **Be unambiguous:** Avoid instructions that could be interpreted in multiple ways.
2. **Enforce a strict JSON schema:** The frontend and backend depend on predictable keys and types.
3. **Constrain actions:** Only allow adaptations that are valid for the given agent type.
4. **Tie reasoning to context:** The model should always consider the event, user profile, and recent history together.

The idea was to make the prompts as predictable as possible. In accessibility systems, consistency often matters more than creativity or flexibility.

4.7.2 Prompt Structure from `sif_config.json`

Each agent (UI, Geometry, Input) has its own prompt in `sif_config.json`. Here's a shortened example from the UI agent:

```

1 {
2   "prompt": "You're the UI suggestion Agent.
3   Analyze user event: {event_json}
4   User profile: {profile_json}
5   Recent history (last 10 events): {history_json}
6   Suggest UI adaptations as JSON in the strict format.
7   Each suggestion must include 'intent', 'reason', and at least a 'value' and
  ↳ 'mode' field. Value must be a reasonable number (e.g., 1.2) with at most
  ↳ one decimal place, and represents a scaling value unless stated otherwise
  ↳ in the metadata (e.g. font size). Target can be 'all' or specific elements."
8 }
```

Listing 4.7: Example UI Agent Prompt

The important elements here are:

- **Role definition:** explicitly stating the agent's focus ("UI suggestion Agent").
- **Context injection:** inserting the current event, profile, and history as JSON strings.
- **Output constraints:** telling the model exactly what keys and value types are allowed.
- **Value rules:** restricting numeric ranges so the model doesn't output absurd sizes.

The complete prompt set, agent-specific instructions, and their runtime configuration are detailed in the codebase of this thesis.

4.7.3 Disjunction Ambiguity in LLM Interpretation

During testing, it was observed that LLMs can misinterpret logical connectors such as **or** and **and**. For example, the instruction:

"The adaptation must include a value or a mode field."

was intended to mean *at least one of these fields is required for this action type*. However, the model sometimes treated this as an exclusive choice (only one allowed) or as fully optional (neither required), even in cases where one was necessary. This can result in incomplete adaptations. This behavior aligns with the well-known inclusive-exclusive disjunction ambiguity described in requirements engineering and computational linguistics, where natural language "or" lacks explicit semantic constraints and is prone to misinterpretation.

For example:

- `{"action": "increase_button_size", "target": "lamp", "value": 1.23}`: requires value but not mode.
- `{"action": "switch_mode", "target": "all", "mode": "voice"}`: requires mode but not value.
- `{"action": "highlight_button_border", "target": "button"}`: valid with neither, as neither field is relevant for this adaptation.

Having both fields is rather uncommon, but not impossible. This should be handled further by the frontend and depends on the developer's implementation choices. Even with this type of prompt engineering, it sometimes still returns a `mode` field when a `value` field would be more appropriate, and vice versa. To reduce misinterpretation, prompts should explicitly state inclusive meaning when a field is required, such as:

“The adaptation must at least include a value and a mode field, depending on the action type.”

This minimises the risk of missing required parameters due to inclusive-exclusive disjunction ambiguity.

4.7.4 Balancing Model Parameters

Even with a well-written prompt, model behavior on the prompt is strongly affected by:

- **Model:** The specific architecture and training of the model can influence its understanding and generation capabilities.
- **Temperature:** how random the outputs are. Lower values (0.2-0.3) are better for consistent JSON.
- **Thinking budget:** how many reasoning steps the model takes before output. Too low, and it may skip checks; too high, and it can slow down or get stuck.
- **Timeout:** waiting time before falling back to mock or rule-based logic.

The Validator Agent especially needs more time and budget because it has a heavier job. This includes more complex reasoning and validation tasks like checking for inconsistencies in the adaptation or checking every adaptation against allowed actions, as well as removing irrelevant adaptations or merging duplicates. In my testing, the validator with a low thinking budget and temperature often hung or exceeded acceptable latency to respond, which is not ideal for real-time adaptations. The goal was to stay under a timeout of max. 30 seconds for the Validator Agent and 15 seconds for the Suggestion Agents.

4.7.5 Avoiding Hallucinations and Bad Values

One common risk with LLM-driven adaptations is hallucination, where the model invents an action, target or value that doesn't exist within the scope of the app. To reduce this:

1. The allowed actions list is always clearly included in the prompt.
2. Targets are validated against the current UI state before applying them as well as by the Validator Agent.
3. A list of focus items is included to guide the model's attention and provide additional context to minimize irrelevant outputs.
4. Prompt clearly states the required fields and their expected values as well as adhering to the allowed actions and JSON contract.

Even so, the validator sometimes has to fix agent mistakes.

For example, if the geometry agent outputs:

```
{"action": "increase_button_size", "target": "button_unlock"}
```

The validator can correct it to:

```
{"action": "increase_button_size", "target": "button_unlock", "value": 1.5}
```

4.7.6 Importance of a Strict JSON Schema

Another critical safeguard against hallucinations and malformed outputs is the use of a strict `response_json_schema` in the LLM API call (in this case Gemini). This schema explicitly defines:

- The **structure** of the output object.
- The **allowed fields**, their types, and descriptions.
- The **required fields** and conditional requirements (e.g., either `value` or `mode` must be present depending on the action type).

By embedding this schema in every LLM request, the backend ensures that:

1. Any response not matching the schema has a higher probability of being rejected before it reaches the frontend.
2. Missing critical parameters (such as `value` for size changes) are caught early.

3. Developers integrating the framework can rely on predictable keys, reducing integration errors.
4. The LLM is gently “steered” towards valid outputs, as many LLM APIs use the schema as a structural hint during generation.

The schema used in this framework is shown below:

```

1 {
2   "type": "object",
3   "properties": {
4     "adaptations": {
5       "type": "array",
6       "items": {
7         "type": "object",
8         "properties": {
9           "action": { "type": "string",
10            "description": "The type of UI adaptation to perform" },
11           "target": { "type": "string",
12            "description": "UI element or component to apply the adaptation to,
13            ↪ 'all' is allowed" },
14           "value": { "type": "number",
15            "description": "Numeric multiplier for size changes (e.g., 1.5 for
16            ↪ 50% larger)" },
17           "mode": { "type": "string",
18            "description": "Interaction or visual mode to switch to (e.g.,
19            ↪ 'voice')" },
20           "reason": { "type": "string",
21            "description": "Why this adaptation was suggested" },
22           "intent": { "type": "string",
23            "description": "Inferred user intent based on the event" }
24         },
25         "required": ["action", "target", "reason", "intent"],
26         "oneOf": [
27           { "required": ["value"] },
28           { "required": ["mode"] }
29         ]
30       },
31     },
32   },
33   "required": ["adaptations"]
34 }
```

Listing 4.8: SIF LLM Output JSON Schema

This schema directly supports accessibility and safety goals: by preventing incomplete or semantically invalid adaptations from being applied, it reduces the risk of unpredictable UI behavior. Combined with the Validator Agent, it forms a two-layer defence. Furthermore, schema validation stops malformed data at the source, while semantic validation ensures that even structurally valid adaptations are contextually appropriate.

4.8 Performance and Evaluation Metrics for AI Logic

This section measures the end-to-end behavior of the SIF and MA-SIF pipelines: *Flutter client* → *Web-Socket* → *FastAPI* → *LLM agents* → *validator* → *client*. The backend processes *one event at a time* (sequential), so the reported latencies reflect a full round-trip with model inference on the critical path. All runs used the same user identifier (`user_seq`). On the very first event, no profile existed, so the server

created a default profile; subsequent events appended to the profile’s `interaction_history` (capped at 10). Qualitatively, later responses carried slightly richer rationales, consistent with the growing history window.

Method

Two lightweight probes were run:

1. **WebSocket round-trip**: a minimal ping-pong to capture end-to-end transport + processing overhead.
2. **Deterministic event suite**: rotation over `tap` → `miss_tap` → `voice` → `gesture`, repeated over a short sequence.

Each response was classified as:

- `validated_by_validator` (final list produced by the validator),
- `combined_agent_suggestions` (validator failed; raw agent outputs returned),
- `mock_rule_fallback` (all agents failed; rule engine response).

The strict JSON output schema used in the backend was enforced to compute “schema-valid %”. To avoid artificial timeouts during heavier MA-SIF runs, client keep-alive pings were disabled. Run sizes were intentionally small (6-10 events) due to API free-tier limits; the goal is to surface architectural trends, not saturate the service.

Configurations

Table 4.1: Exact agent settings for the three measured configurations.

Config	Agent	Model	Temperature	Thinking Budget	Timeout (s)
SIF (single agent)	SIF agent	gemini-2.5-flash	0.2	dynamic (-1)	api-default
MA-SIF (balanced)	UI agent	gemini-2.5-flash-lite	0.2	0	15
	Geometry agent	gemini-2.5-flash-lite	0.2	0	15
	Input agent	gemini-2.5-flash-lite	0.2	0	15
	Validator	gemini-2.5-flash	0.3	dynamic (-1)	30
MA-SIF (heavy)	UI agent	gemini-2.5-flash	0.2	2048	30
	Geometry agent	gemini-2.5-flash	0.2	2048	30
	Input agent	gemini-2.5-flash	0.2	2048	30
	Validator	gemini-2.5-flash	0.3	2048	30

Results

Table 4.2: Latency and correctness across configurations (sequential backend; `user_seq` with default profile and growing history).

Config	WS p50 (ms)	WS p90 (ms)	WS max (ms)	Suite p50 (ms)	Suite p90 (ms)	Suite max (ms)	Schema-valid (%)	n
SIF (single agent)	8240.99	9913.63	10649.21	8261.86	12034.12	12914.24	40.0	10
MA-SIF (balanced)	9969.99	11151.80	11164.61	10988.27	14659.71	15234.09	70.0	10
MA-SIF (heavy)	26875.99	27613.81	32009.47	27229.24	30443.48	31347.16	100.0	10

Legend. WS p50/p90/max: WebSocket round-trip latency percentiles and maximum (ms). Suite p50/p90/max: end-to-end latency percentiles and maximum for the deterministic event suite (ms). Schema-valid: share of responses that passed JSON schema validation (%). *n*: number of events.

Across all runs 0% `combined_agent_suggestions` and 0% `mock_rule_fallback` was observed: the validator returned a final list for every event. In the heavy configuration, some individual agent calls failed (API-side), yet the validator still produced a valid final list, so external behavior stayed stable. (Table 4.2)

Note: in earlier exploratory tests, fallbacks mainly appeared when the free-tier quota was hit or the validator timed out, which did not occur here. Furthermore, during this evaluation it was observed that the validator agent in the MA-SIF heavy config accepted more suggestions from the agents than in the other more balanced config.

Interpretation

The following patterns can be observed:

- Moving from **SIF (single agent)** to **MA-SIF (balanced)** increases median latency from ~ 8.3 s to ~ 11.0 s, but schema validity jumps from **40%** to **70%**. Even with zero thinking budget on the specialist agents, the *validator* is pulling its weight.
- **MA-SIF (heavy)** pushes median latency to ~ 27.2 s, but schema validity reaches **100%**, and the *reason/intent* fields become noticeably richer and more proactive.
- The short history window helps later events rationales get slightly better as context accumulates, although the biggest quality gain clearly comes from the validator’s (higher) thinking budget.

Practical read: For accessibility-centric flows, a sensible pattern is:

1. **Immediate, local rules** for fast feedback (e.g., enlarge target, boost contrast).
2. **MA-SIF in the background** for the smarter, profile-aware follow-ups (e.g., switch to voice, simplify layout).

This preserves responsiveness while still getting the benefits of multi-agent reasoning.

Limitations

The runs were kept small (6-10 events) due to free-tier limits, so statistical power is modest—though the differences are large and consistent. The backend is sequential, so times reflect worst-case per-event latency with LLM on the critical path. The event suite is deterministic; other UI states can nudge absolute times, but the overall trade-off (latency \leftrightarrow schema quality) is expected to hold.

Implications for the Framework

Overall, the numbers back the architecture: MA-SIF + a capable validator improves output predictability and schema adherence; latency is tunable via thinking budgets, model choices and timeouts.

Recommended deployment profiles:

- **Tiered response:** apply *deterministic* local adaptations instantly; let MA-SIF deliver deeper, profile-aware changes a few seconds later.
- **Balanced default:** *gemini-flash-lite* for UI/geometry/input agents (`thinking_budget=0`, `timeout=15s`); *gemini-flash* validator with dynamic thinking and `timeout=30s`.
- **Heavy mode (opt-in):** enable larger validator thinking budgets for flows that demand airtight structure and richer rationale (e.g., clinical or regulated settings). Expect $\sim 2\text{--}3\times$ latency vs. balanced.

Guidelines for Model Optimization:

- **Enhanced Accuracy and Reduced Schema Errors:** Prioritize increasing the validator thinking budget for optimal cost-effectiveness.
- **Reduced Latency:** Utilize *gemini-flash-lite* agents with zero thinking time and set a validator timeout cap.
- **Comprehensive Explanations:** Increment validator thinking budget first, followed by agent budget, while maintaining low temperatures (0.2-0.3).

Failure handling Even with some agent API errors, the validator still made valid final lists. Returning other agent adaptations when the validator fails will provide a great fallback while the mock rule fallback is still the best backup for quota or timeout problems on all agents, although both are barely used in normal situations.

Conclusion Use MA-SIF for better quality and consistency, but only deal with the slower latency when it's a necessity. Adding local rules will help keep the UI smooth for accessibility needs. As profiles and histories get bigger, there will be small but steady improvements in the reasoning quality over time, although the validator's budget is still the main way to control accuracy. Using the gemini-pro model could potentially improve adaptations even further in important scenarios.

4.9 Limitations and Solutions of LLM Integration

Large Language Models make Smart Intent Fusion far more capable than any static rule set, but they also introduce new dependencies, performance constraints, and reliability issues. In this thesis, all reasoning was powered by the Gemini API, and while this enabled rapid development, it also shaped both the strengths and weaknesses of the final system.

4.9.1 LLM selection

Gemini was chosen as the sole LLM provider for this framework for a mix of practical and technical reasons. The generous free tier and large input/output token limits allowed for frequent iteration without cost becoming a limiting factor. The API provided access to both smaller, faster models (used for the UI, Geometry, and Input agents) and larger, reasoning-focused models such as `gemini-2.5-flash` or even `pro` (used for the Validator Agent). This made it possible to optimize speed where possible and allocate more resources to roles that needed deeper analysis.

Gemini's handling of structured JSON output was also an advantage, as Smart Intent Fusion depends on predictable schema-compliant responses. While this project did not directly benchmark other models like GPT-5 or Grok, partly to keep the system stable during development and partly because Gemini's free tier already covered the thesis's usage without cost. Stability was important for building and testing MA-SIF without constantly adjusting prompts and entire agents for different model behaviors. However, it also means that all testing and performance observations in this chapter are specific to Gemini's runtime behavior, and the results may differ if another LLM were used.

4.9.2 Reliability and Latency Constraints

A constant challenge was balancing response quality with the speed needed for near real-time accessibility. Smaller "lite" models returned in fractions of a second and were ideal for the three suggestion agents. The Validator Agent, however, required the larger `gemini-2.5-flash` model to perform more complex and reliable checks across multiple agent outputs. The trade-off was that validation could sometimes become the slowest part of the pipeline, especially with its increased timeout (30 seconds) and dynamic thinking budget. Network latency or temporary API slowdowns in the agents could lead to fewer suggestions being returned for a given event. In those cases, partial results were still applied rather than waiting or retrying for a full set.

4.9.3 Hallucinations and Invalid Output

Even with strict prompts and explicit allowed-action lists, hallucinations still appeared in the output. These took the form of actions outside the approved set, targeting elements that did not exist, or producing unreasonable scaling values (for example, value: 10). The Validator Agent was able to remove or correct most of these before they reached the frontend, but this came at the cost of extra processing time and complexity. Without validation, such outputs could have caused visual glitches or broken layouts, especially in geometry-related adaptations. The Input Adapter Layer or the Frontend itself, should as a fallback also have extra validation checks.

4.9.4 Token Limits and Context Size

Gemini’s token allowance was one of the main reasons for choosing it, but token limits were still a factor. Each agent’s prompt included the current event, the user profile, and the last few events from history, all in JSON format. In cases where the history was long or metadata verbose, this could approach the model’s input size limit. The solution was to truncate history in those cases, ensuring that the event and profile data always took priority, even if it meant losing some recent context. Furthermore the token limit per minute and per day also limited the number of requests that could be made, which is why the framework was designed to handle partial failures gracefully and still return useful adaptations even if some agents timed out or exceeded limits.

4.9.5 Validator Complexity

The Validator Agent is both the most important and the most resource-intensive part of MA-SIF. It merges outputs from multiple agents, removes duplicates, corrects invalid values, resolves conflicting suggestions, and ensures the final adaptations list passes schema validation. This workload made it prone to timeouts when handling a large number of adaptations at once. Increasing the thinking budget and timeout reduced these failures but also increased total response time, creating a constant balance between reliability and speed.

4.9.6 Dependency on External APIs

Finally, using an external LLM API means the framework is dependent on network availability and the stability of the provider. If the Gemini API is unavailable or returns errors, the system falls back to the rule-based logic. While this ensures baseline functionality, it also removes the more context-aware reasoning that makes Smart Intent Fusion valuable. In a production setting, this could be mitigated with on-device models or by integrating multiple LLM providers as backups.

4.10 Future Directions for AI-Driven Adaptation

While Smart Intent Fusion in its current form is functional and effective for the scenarios tested in this thesis, it is still a first iteration of what an AI-driven, multimodal UI adaptation system could be. The underlying architecture, especially the multi-agent design and the strict JSON-based API contract, was deliberately built with future extensions in mind. Several directions could significantly expand its capabilities and make it more adaptive and autonomous.

One natural evolution is the introduction of visual context through image, vision-based or even at runtime UI analyzer models. Currently, SIF relies on structured metadata from the frontend to understand the UI state. In future versions, a lightweight computer vision model could take periodic screenshots of the interface and produce a semantic map of UI elements, their sizes, positions, and visual contrasts. This map could then be passed to the LLM alongside the existing event, profile, and history data. The advantage of this approach is that the AI would be reasoning over actual UI layouts, rather than relying on the frontend to describe them accurately which can cause misinterpretations. This opens the door to truly context-aware adaptations. For example, increasing the size of the smallest actionable element on a crowded screen, even if it hasn’t yet caused an interaction error or accurately reposition elements closer to the user’s focus.

Another promising direction is integrating user feedback into the adaptation loop. At present, SIF updates the user profile implicitly, based on interaction patterns or the user itself. A future version could prompt the user after significant adaptations with a quick, accessible feedback mechanism (“Was this change helpful?”). This feedback could be stored alongside interaction history and used by the LLM to adjust its decision-making over time.

There is also scope for dynamic, UI-level code changes driven directly by the LLM. Currently, SIF works within a fixed set of allowed actions and values prone to some hallucinations. This could be expanded so that the LLM can modify layout constraints, create new UI elements, or reorganise screens entirely with safeguards in place to prevent breaking the interface. The LLM could use its own “hallucinations” to provide creative solutions for layout issues or user interactions. This would take SIF from an adaptation system to a full UI changing layer, capable of designing new interactions on demand.

Finally, the framework could explore multi-model, multi-provider reasoning and threading. At present, all reasoning is performed by Gemini and run sequentially, which simplifies development but limits the diversity and speed of outputs. Future versions could run agents across different LLM providers and threads or even mix LLMs with specialised non-language models (e.g., reinforcement learning agents for adaptation strategies), with the Validator Agent controlling and validating the final output. This would provide extra resilience against provider outages and allow different models to play to their strengths, as well as a strong speed improvement when asynchronous processing is implemented. Every agent could potentially run in its own thread or process and joined by the validator, allowing for true parallelism and faster overall response times.

4.11 Chapter Summary

This chapter presented Smart Intent Fusion (SIF) as the core reasoning layer of the adaptive framework. SIF combines multimodal inputs, user profiles, and recent interaction history to deliver personalised, context-aware UI adaptations. A hybrid approach of rule-based logic and LLM-driven reasoning ensures both reliability and flexibility, allowing essential accessibility features to work instantly while enabling more complex, context-sensitive changes.

The multi-agent architecture (MA-SIF) was introduced, with specialised agents for UI, geometry, and input adaptations, and a Validator Agent responsible for merging and cleaning outputs. This design improves reliability, supports partial fallbacks, and can be dynamically reconfigured via `sif_config.json`. Prompt engineering emerged as a critical factor in ensuring valid, schema-compliant LLM output, with careful wording required to avoid logical misinterpretations.

Limitations of LLM integration including latency, occasional hallucinations, and reliance on a single provider, were mitigated through strict validation and fallback mechanisms. Finally, the chapter outlined future directions for SIF, such as adding visual UI context, integrating user feedback, enabling deeper UI changes, and exploring on-device AI models.

Overall, Smart Intent Fusion was presented not just as an algorithm, but as a modular, extensible smart reasoning layer designed to work across platforms, adapt to different users, and remain robust in the face of LLM unpredictability. It is the component that turns multimodal input into meaningful, personalised adaptations for reshaping the interface to fit the user.

Chapter 5

An Adaptive Multimodal GUI Framework using LLMs

5.1 Introduction to an Adaptive Smart Home Controller

The Adaptive Smart Home Controller ¹ is the practical proof-of-concept used to implement and validate the multimodal AI-driven GUI framework presented in this thesis. It serves as a concrete example of how the framework’s concepts, introduced in Chapter 3, can be applied in a real, interactive application. The Smart Home Controller simulates the control of typical household devices such as lights, thermostats, and door locks. While the devices themselves are virtual, the process of capturing inputs, interpreting them through the backend with user profiles and history, and applying adaptive changes to the interface is authentic and functionally representative of a real deployment.

The choice for a smart home context was deliberate. It offers a relatable and real-life use case structured set of interaction scenarios that cover a range of UI components, buttons, sliders, text elements which are central to accessibility-focused adaptations. Furthermore, it reflects real-world situations where users may have diverse abilities and input preferences. For example, a motor-impaired user might need larger buttons to avoid frequent miss-taps, while a visually impaired user may benefit from higher-contrast modes and enlarged text. By embedding these scenarios into a single, unified application, the Smart Home Controller provides a manageable but representative testbed for the framework’s adaptive capabilities.

The system adheres to the three-layer architecture established in earlier chapters. The frontend layer, implemented in Flutter (`adaptive_ui_app.dart`), renders the interface, captures user interactions, and applies adaptation instructions as they are received from the backend. The Input Adapter Layer (`adaptive_ui_adapter.dart`) acts as a middleware component, converting raw inputs from multiple modalities into the framework’s JSON-based event format and ensuring user profiles are retrieved or updated before events are transmitted. The backend layer (`backend.py`), built with FastAPI, Gemini API and MongoDB, implements the Smart Intent Fusion (SIF) process. This includes both deterministic, rule-based adaptations and multi-agent LLM-driven reasoning (MA-SIF), combining event data, user profiles, and interaction histories to produce targeted adaptation actions.

It is important to note that this first iteration does not yet incorporate every capability described in the framework’s long-term vision. Certain modalities, such as gesture recognition, are currently simulated through mock events to keep the implementation focused on the adaptation pipeline rather than input hardware integration. Throughout this chapter, each component is discussed in detail, with a clear distinction made between fully implemented functionality, simulated elements, and features that remain as future work.

¹The codebase containing implementation and evaluation data of this thesis can be found at: <https://github.com/YarneD-1952226/A-Multimodal-AI-Driven-GUI-Framework-for-Dynamic-User-Adaptation>

5.2 Development Environment

The development environment for the Adaptive Smart Home Controller was chosen to support rapid prototyping, cross-platform deployment, and straightforward integration with the multimodal adaptation framework described in earlier chapters. It needed to provide a fast feedback loop during implementation, flexibility in UI design, and robust backend capabilities to support real-time Smart Intent Fusion. The final setup reflects a balance between practical constraints such as available hardware, time and technical requirements namely WebSocket support, database persistence, and LLM integration.

Flutter was selected for the frontend because of its ability to produce visually consistent applications across desktop, mobile, and web from a single codebase. The framework's reactive UI model and composable widget system made it straightforward to create adaptive components whose properties such as size, color, and layout can be updated dynamically in response to backend instructions. Flutter's hot reload feature also significantly reduced iteration time, which proved essential for testing the frequent, small adjustments needed when refining adaptation behaviors.

The backend is implemented in Python using FastAPI, chosen for its simplicity, asynchronous request handling, and native WebSocket support. FastAPI's lightweight structure allowed the project to keep the adaptation pipeline transparent and easily modifiable, while still offering the performance needed for real-time interactions. MongoDB was selected as the database because its document-based structure aligns directly with the JSON contracts used throughout the framework. It stores user profiles, interaction histories, and adaptation logs without the need for complex schema migrations, making it well-suited for iterative development.

Development and primary testing took place on macOS 15.6, with additional runs on Windows 11 to confirm cross-platform compatibility. Both environments used Visual Studio Code with the Flutter and Dart plugins for frontend work, and Python tooling for backend development. This combination made it possible to run and debug both layers side-by-side, with the frontend communicating directly to a locally hosted backend via WebSocket and HTTP.

For clarity, the main environment components were as follows:

- **Operating System:** macOS 15.6 (primary), Windows 11 (secondary testing)
- **Frontend Framework:** Flutter SDK 3.7.0 or higher
- **Programming Languages:** Dart 2.19.0+ (frontend), Python 3.9+ (backend)
- **Backend Framework:** FastAPI with Uvicorn ASGI server
- **Database:** MongoDB 6.0+ for persistent profile and interaction history storage
- **IDE:** Visual Studio Code with relevant Flutter/Dart and Python extensions
- **Communication:** WebSocket for real-time adaptation updates, HTTP for profile management and batch operations
- **Version control:** Git, with the repository hosted on GitHub for collaboration and version tracking.

While this configuration is sufficient for the current implementation, it is also designed to be portable. The backend can be deployed to cloud environments without modification, and the frontend can target iOS, Android, or desktop platforms simply by recompiling or adjusting the build configuration. This flexibility ensures that the same codebase can serve as both a research prototype and a potential foundation for future production-ready systems.

5.3 Frontend (Flutter): Adaptive Smart Home Controller

The frontend of the Adaptive Smart Home Controller is implemented in Flutter and serves as the primary user-facing component of the framework. Its role is to render the interface, capture user interactions across multiple modalities, and apply adaptation actions received from the backend in real-time. Although the backend is responsible for reasoning about what adaptations to make, the frontend is where these adaptations become visible to the user and directly influence usability.

5.3.1 Responsibilities & Data Flow

The Flutter frontend (`AdaptiveSmartHomeApp` → `SmartHomeScreen`) renders the device cards, captures user interactions (live touch; voice/gesture mocked), and applies backend-issued adaptations in real-time via a WebSocket callback. The `AdaptiveUIAdapter` is constructed with a user ID and a callback; when the backend responds with a list of `UIAdaptation` items, the frontend updates reactive state maps and triggers lightweight animations. (Figure 5.1)

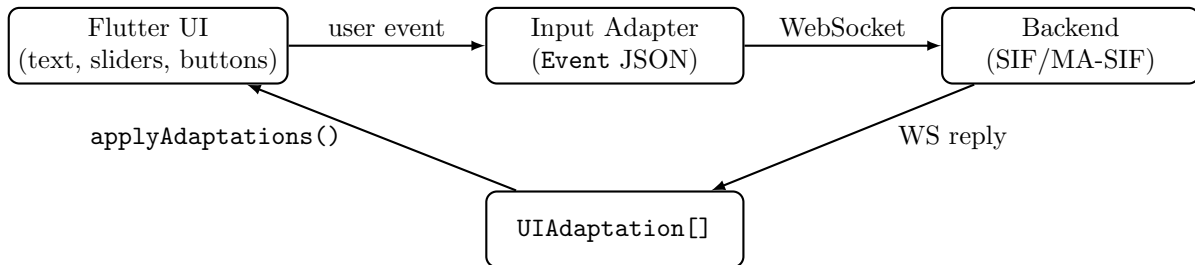


Figure 5.1: Frontend data flow: events to adapter; adaptations back to UI.

5.3.2 UI Composition & State Model

The UI is a scrollable list of device cards with minimalist controls each representing a smart home device such as a lamp, thermostat or a door. These cards contain core interactive elements, including buttons for on/off actions, sliders for settings such as temperature, and labels for contextual information. (Figure 5.2)

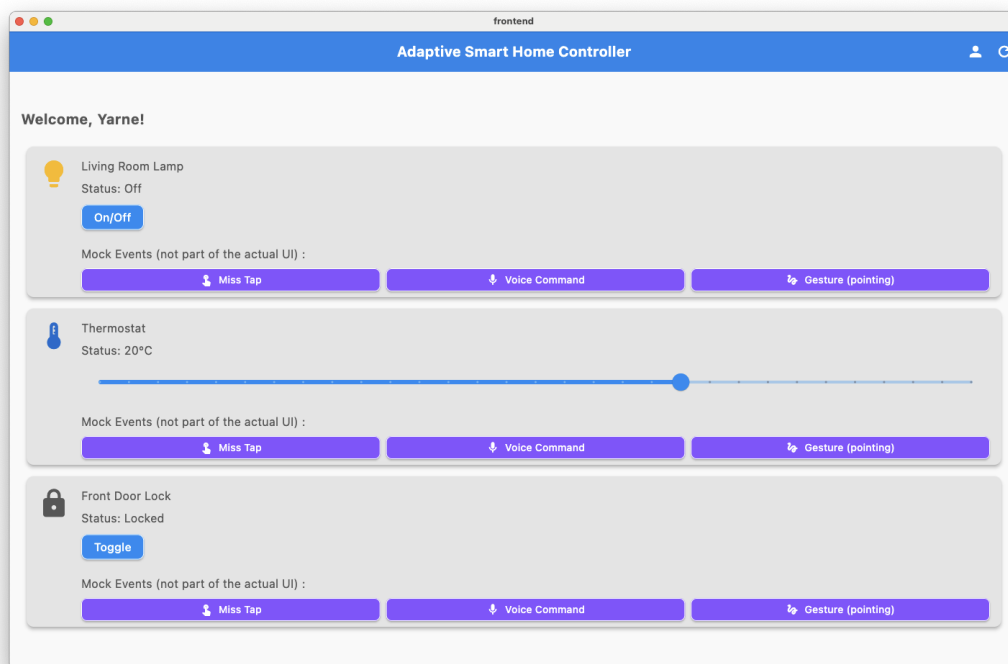


Figure 5.2: Smart Home Controller UI: scrollable device cards (Lamp, Thermostat, Lock) with minimalist controls and the mock event row (*Miss Tap*, *Voice Command*, *Gesture*).

The UI is intentionally minimalist, focusing on clarity and accessibility. Each device card (Lamp, Thermostat, Lock) presents its controls in a clean, uncluttered layout: large buttons for toggling states, sliders for adjustments, and concise status labels. The design prioritizes high visual contrast and generous spacing, making it easier for users with motor or visual impairments to interact. Adaptations such as increased button size, font scaling, or contrast changes are immediately reflected in the UI, with smooth animations

to reinforce feedback. A persistent mock event row below each card allows users (or testers) to simulate miss-taps, voice commands, or gestures, demonstrating how the interface responds and adapts in real-time. This modular structure ensures that accessibility enhancements are both visible and intuitive, supporting a wide range of user needs without overwhelming the interface.

5.3.3 Event & Adaptation Contract

Event (frontend→backend). Events carry modality and UI context:

```

1 {
2   "eventType": "miss_tap", // e.g., touch, voice, gesture, miss_tap, ...
3   "source": "touch", // modality
4   "target": "lamp", // lamp | thermostat | lock
5   "metadata": { "UI_element": "button" } // or slider, + command/gesture fields
6 }
```

Listing 5.1: Event structure

Note. The Dart Event uses camelCase fields (e.g., `eventType`, `targetElement`); the adapter serializes to `snake_case` for the JSON contract.

This event will later be enriched with other required fields in the Input Adapter Layer, such as `timestamp` and `user_id`.

Adaptation (backend→frontend). The backend replies with a list of atomic actions:

```

1 [
2   { "action": "increase_button_size", "target": "lamp", "value": 1.2 },
3   { "action": "increase_font_size", "target": "all", "value": 1.15 },
4   { "action": "increase_contrast", "target": "all" },
5   { "action": "show_tooltip", "target": "lamp", "value": "Try a longer press" }
6 ]
```

Listing 5.2: Adaptation example actions

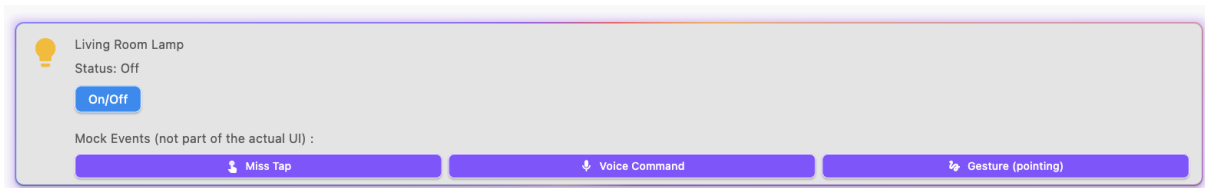
Adaptations received from the backend after being processed by the Input Adapter Layer are applied immediately using Flutter’s reactive state management. This is handled in the `applyAdaptations(...)` callback function. For example, a `increase_button_size` action triggers an `AnimatedScale` widget update, enlarging the targeted button over a short animation to make the change more noticeable without disrupting the user’s flow. Similarly, a `increase_contrast` action adjusts the application’s color scheme by updating theme parameters, while text-related adaptations update font sizes dynamically. Furthermore, when these adaptations reach the frontend, a brief highlight at the top of the application, indicates the changes being applied. This direct mapping between adaptation actions and Flutter widget properties allows the frontend to respond flexibly to a wide range of changes without requiring hardcoded layouts. See Table 5.1 for details.

Table 5.1: Adaptation-to-widget mapping (from `applyAdaptations`).

Action	Target	Effect in UI
<code>increase_button_size</code>	device or all	Scales button via <code>buttonScales[target]</code> (animated).
<code>increase_button_border</code>	device or all	Doubles border thickness in <code>elementBorders</code> .
<code>increase_slider_size</code>	device or all	Multiplies slider scale in <code>sliderSizes</code> .
<code>increase_font_size</code>	(global)	Multiplies all entries in <code>fontSizes</code> .
<code>increase_contrast</code>	(global)	Switches to high-contrast theme via <code>_switchToHighContrastTheme()</code> .
<code>adjust_spacing</code>	device or all	Scales spacing in <code>elementSpacing</code> .
<code>show_tooltip</code>	device	Floating <code>SnackBar</code> with helper text.
<code>switch_mode</code>	(global)	Records navigation-mode change (placeholder for future).
<code>trigger_button</code>	lamp/lock	Toggles <code>deviceStatuses</code> (<i>On/Off</i> , <i>Locked/Unlocked</i>).
<code>simplify_layout</code>	(global)	Sets <code>simplifiedLayout=true</code> to reduce clutter.

5.3.4 Responsiveness & Feedback

While awaiting adaptations, the active card displays a rotating gradient border implemented via the `_AnimatedGlowBorder` widget. This subtle animation serves as a visual indicator of system latency, signaling to users that their input is being processed and an adaptation is forthcoming. By using a non-intrusive glowing effect, the interface maintains user engagement without causing distraction or interrupting ongoing interactions (see Figure 5.3). This approach is particularly beneficial for accessibility, as it provides clear feedback for users who may rely on visual cues to understand system status.

**Figure 5.3:** Loading indicator: rotating gradient card border while waiting for backend.

5.3.5 Profile Bootstrap & Editing

On first run, the frontend checks if a profile exists; if not, it creates a default profile and persists it via the adapter. A bottom-sheet editor (`ProfileEditorSheet`) allows the ability to edit accessibility flags and UI preferences (font size, button size) of the user, see figure 5.4, then calls `updateProfile(...)`.

```

1  if widget is initialized then
2      adapter := AdaptiveUIAdapter(user_id, adaptation_callback)
3      if adapter.checkProfile(user_id) then
4          profile_data := adapter.getProfile(user_id)
5          user_profile := parseProfile(profile_data)
6      else
7          user_profile := UserProfile(
8              accessibility_needs=default,
9              input_preferences=default,
10             ui_preferences=default,
11             interaction_history=[]
12         )
13         adapter.updateProfile(user_profile)
14  end if

```

Listing 5.3: Profile initialization logic

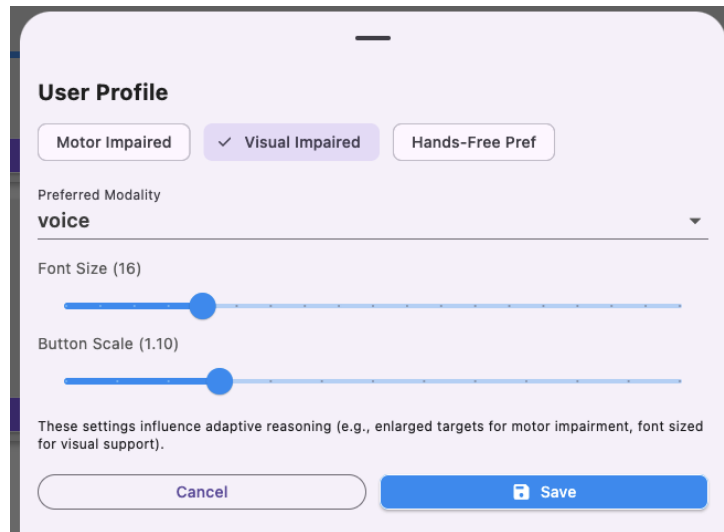


Figure 5.4: Profile editor for accessibility needs and UI preferences.

5.3.6 Testing Harness: Mock Events & Input Capture

To validate the end-to-end pipeline without physical input hardware, a mock panel triggers *miss tap*, *voice command*, and *gesture* events per device. Each button constructs an **Event** with modality, target, and UI context and sends it via `adapter.sendEvent(...)`. This keeps the architecture ready for real speech/gesture modules later, while staying reproducible for the study.

When a user presses on an mockup event in the test panel underneath each device, the frontend captures the event along with metadata such as the element’s type for UI context, the interaction type, and any other relevant parameters. This data is sent through the Input Adapter Layer, which standardizes the event into the framework’s JSON contract before forwarding it to the backend via WebSocket. The choice of WebSocket ensures that adaptations can be returned and applied in near real-time, a critical requirement for maintaining a smooth user experience in accessibility scenarios. The power of the framework lies in the ability to set the `event_type`, `source`, and others, dynamically based on user interactions, since the backend will use LLM-reasoning to interpret these events in context. While the fields are bound by the JSON contract to standardize events (for modularity and consistency), the actual values can be highly contextual and dynamic.

Input Capture: The frontend layer captures raw inputs from multiple modalities, including touch, voice, and gestures. Each modality is handled through specific Flutter widgets and event listeners:

- **Touch:** Taps or miss-taps on buttons/sliders, detected via Flutter’s `GestureDetector` or `onPressed` callbacks. Miss-taps are mocked but extensible to hover or real misstep detection like `MouseRegion` or `Listener` widgets, to create a “detection bounding box”.
- **Voice:** Commands like “Turn on lamp” or “Unlock door”, mocked for this thesis but extensible to libraries like `speech_to_text` or Web Speech API.
- **Gestures:** Hand movements (e.g., point, swipe) via mock events, with future support for MediaPipe.

5.3.7 Summary

Overall, the Adaptive Smart Home Controller (`adaptive_ui_app.dart`) demonstrates how the frontend layer of the framework can be implemented in a way that is both platform-independent and responsive to dynamic adaptation instructions. By separating UI rendering from adaptation logic and using the

Input Adapter Layer as a bridge, the application remains modular, making it easier to extend or replace individual components without affecting the overall system.

5.4 Input Adapter (Dart): Transport, Serialization & Adaptation Callback

The adapter bridges the Flutter frontend and the SIF backend by (i) serializing internal **Event** objects into the JSON event contract, (ii) handling low-latency transport via WebSocket for adaptations, and (iii) managing user profiles over HTTP. This section focuses on the concrete implementation in `adaptive_ui_adapter.dart`: the class surface, serialization choices, transport lifecycle, and the `onAdaptations` callback function into the frontend. See Chapter 3 for the high-level event schema and field definitions.

5.4.1 Class Overview

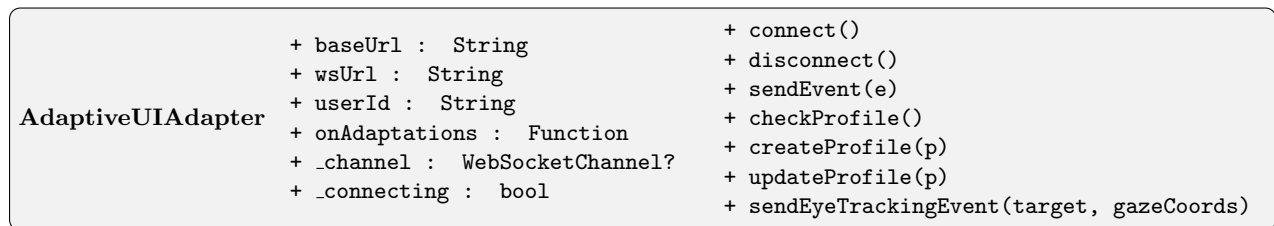


Figure 5.5: Simplified UML-like diagram for `AdaptiveUIAdapter`.

In the current implementation (Figure 5.5), the adapter intercepts mock events generated by the frontend, such as missed button presses, missed slider adjustments, or simulated voice commands. Each event is enriched with metadata, including the user’s identifier, a timestamp, the type of interaction, and any target element references. The adapter then converts this information into the JSON event contract, further defined in Chapter 3, which serves as the standard interface between the frontend and backend. This contract includes fields for event type, source modality, target element, coordinates if applicable, confidence level, time stamp, `user_id` and additional metadata such as the spoken command in the case of voice input.

5.4.2 Internal Representations of Event, Adaptation and User Profiles

The adapter defines core Dart classes for representing interaction events, user profiles, and adaptation actions. The **Event** class encapsulates all relevant fields such as event type, source modality, target element, etc., ensuring every interaction is consistently structured and easily serializable to the backend’s JSON contract (see figure 5.6 for schema details). Similarly, **UIAdaptation** models adaptation instructions received from the backend, while **UserProfile** maintains accessibility needs, preferences, and recent history. These unified data structures guarantee reliable, type-safe communication between frontend and backend, simplify integration, and support extensibility for future modalities or adaptation types.

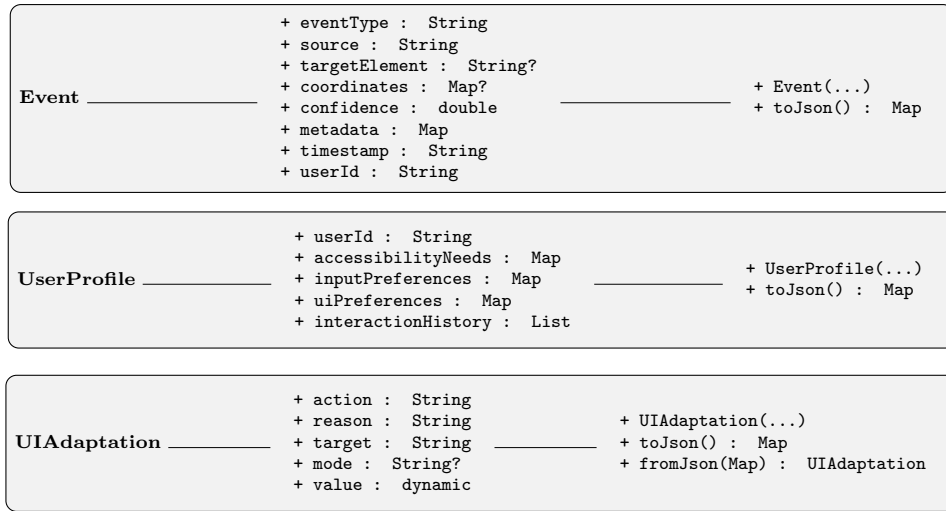


Figure 5.6: High-level class diagram for Input Adapter Layer data structures.

5.4.3 Transport & Profile Management

Before forwarding an event, the adapter queries the backend via HTTP to check whether a profile exists for the given user. If no profile is found, it prompts the frontend to initiate a profile creation request, using default parameters or pre-filled accessibility preferences where available. This mechanism prevents situations where adaptation requests are processed without the necessary user context, which could lead to ineffective or even counterproductive UI changes.

Communication with the backend is handled primarily through WebSocket for low-latency adaptation feedback. HTTP requests are used for profile management, batch operations, and other non-real-time interactions. This division ensures that profile updates and administrative tasks do not interfere with the responsiveness of live adaptations, see table 5.2 for details.

Table 5.2: Adapter endpoints used by `adaptive_ui_adapter.dart`.

Purpose	Method/Proto	Path
Realtime adaptations	WebSocket	/ws/adapt
Check profile exists	GET	/profile/{user_id}
Create profile	POST	/profile
Update profile	POST	/profile

WebSocket Lifecycle & Error Handling:

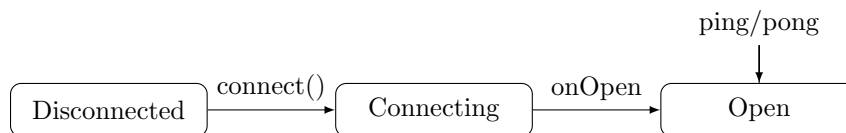


Figure 5.7: Adapter WebSocket lifecycle with reconnect backoff.

In `adaptive_ui_adapter.dart`:

- On **open**: subscribe to messages and (optionally) flush any queued events.
- On **message**: parse JSON array/object, convert to `UIAdaptation` list, invoke `onAdaptations`.
- On **error/done**: close channel.
- **Keep-alive via ping/pong**: Keep-alive via ping/pong is handled at the protocol level; the adapter does not send explicit pings.

Sending Events:

When the frontend is ready to send an event, it calls the `sendEvent` function from the adapter, which enriches the event data with metadata (e.g., `user_id`, `timestamp`), serializes it into JSON, and sends it over the WebSocket connection.

```

1 void sendEvent(Event eventData) {
2     eventData.timestamp = DateTime.now().toIso8601String();
3     eventData.userId = userId;
4     channel!.sink.add(jsonEncode(eventData.toJson()));
5 }

```

Listing 5.4: `sendEvent`: serialize and send over WS**Receiving Adaptations:**

When the adapter receives adaptations from the backend in the WebSocket channel, the adapter processes the JSON and converts them in a `UIAdaptation` list. This list is then passed to the `onAdaptations` frontend callback function. This function is responsible for updating the UI based on the received adaptation data, see Section 5.3.3 for more details on this process.

Profile Management:

Profile management is mostly done by the frontend, with the Input Adapter Layer providing all the necessary functions and ensuring that every user interaction is contextualized with the correct accessibility preferences and history. Before sending any event to the backend, the frontend (as described earlier) checks whether a profile exists for the current user by calling `checkProfile()`. If no profile is found, it creates one using `createProfile(UserProfile p)`, which serializes the profile data and posts it to the backend. Updates to the profile, such as changes in accessibility needs or UI preferences, are handled by `updateProfile(UserProfile p)`, which issues a POST request with the updated profile JSON. These functions guarantee that the backend always receives events enriched with up-to-date user context, supporting personalised adaptation and continuous learning.

5.4.4 Extensibility Example

As proof-of-concept for the extensibility of the framework, a gaze event hook was implemented that incorporates eye-tracking data into an event. This demonstrates how easily new input modalities can be integrated into the existing architecture.

```

1 void sendEyeTrackingEvent(String target, Map<String,double> gazeCoords) {
2     sendEvent(Event(
3         eventType: 'eye_tracking',
4         source: 'gaze',
5         targetElement: target,
6         coordinates: gazeCoords,
7         confidence: 1.0,
8     ));
9 }

```

Listing 5.5: Gaze event hook (testing harness)

Note: This can either be done in the frontend or directly in the adapter, depending on the specific requirements of the application. For demonstration purposes, the gaze event hook was implemented in the adapter, while the frontend would be a more logical place for it in a production setting.

5.4.5 Summary

The adapter is fully wired for touch events and for mocked voice/gesture events triggered from the frontend's test panel, enabling end-to-end validation without external hardware. The adapter treats

mocked and real events identically, so integrating actual speech/gesture modules later is a frontend-only change. By isolating serialization, profile checks, and transport here, the frontend stays focused on rendering, and the backend receives consistent inputs. The same adapter design can be reused in SwiftUI or other platforms with minimal changes to the method surface and JSON contract.

5.5 SIF Backend Layer: Implementation of Adaptation Logic

The backend implements the decision-making core of the framework. It receives standardized events from the adapter, fuses them with user profiles and recent interaction history, and returns concrete adaptation actions for the frontend to apply. The service is written in Python using FastAPI, with Uvicorn for serving requests and MongoDB for persistent storage of profiles and logs. WebSocket is used for low-latency, bidirectional communication during interaction, HTTP is used for profile management and auxiliary endpoints.

5.5.1 Webserver layout and endpoints

The application exposes a WebSocket endpoint, `/ws/adapt`, that accepts JSON events matching the contract introduced earlier. Each message is parsed into a Pydantic `Event` model, the user profile is loaded from MongoDB, and the event, profile, and short history window are passed to the Smart Intent Fusion routine. The resulting adaptation list is returned on the same socket, allowing the frontend to update the UI immediately. For non-interactive operations the backend offers `POST /profile` for profile creation, `GET /profile/user_id` for retrieval, and a small set of diagnostic endpoints such as `GET /full_history` and `GET /modalities`. Profile writes use FastAPI `BackgroundTasks`, which keeps the interaction path responsive while updates are persisted asynchronously.

Startup & Runtime

For local runs, the service is started via a small shell script (`start_backend_processes.sh`) that prepares a venv, checks MongoDB, and runs Uvicorn with explicit WebSocket keepalive settings:

```
1 uvicorn backend:app --reload --host 0.0.0.0 --port 8000 \
2   --ws websockets --ws-ping-interval 60 --ws-ping-timeout 180
```

Listing 5.6: Backend startup (excerpt)

This configuration enables reliable long-lived WS sessions during interactive testing, and eases startup of all the different services.

5.5.2 Data persistence and history management

MongoDB stores user profiles in a `profiles` collection and event-adaptation pairs in a `logs` collection. Profiles are indexed by `user_id` to allow direct lookups during interaction. Each time an event is processed the backend appends a compact JSON representation to the profile's `interaction_history`, capped to a small sliding window. This keeps the prompt context focused on recent behavior while avoiding unnecessary growth in the database. For redundancy, a JSONL file mirrors the log entries during development, which simplifies offline inspection or debugging purposes when the database is reset.

5.5.3 Smart Intent Fusion and MA-SIF

The fusion step supports two paths. A multi-agent configuration, MA-SIF, is the default. It loops over a set of specialised LLM agents, UI, Geometry, and Input, each prompted with the current event, the user profile, and a short history via runtime injection (for more details see Chapter 4), defined in the `sif_config.json`. These agents propose structured adaptations in their domain, for example increasing button size, adjusting spacing, switching input mode, or triggering a button when intent is clear. Their outputs are then passed to a dedicated Validator agent that consolidates, filters, and normalises the suggestions into a final list. The validator removes duplicates, corrects out-of-range values, and ensures that every adaptation conforms to the allowed action set and includes a target and either a value or a mode and more. See Table 5.3 for the current agent configuration.

A single-agent SIF path is also available. It produces a complete adaptation list in one call and is useful when quick iteration is preferable over agent specialisation. Both paths share the same I/O schema, which keeps the frontend indifferent to which reasoning strategy is currently set to active.

Current Agent Configuration (MA-SIF balanced) (from `sif_config.json`)

Agent	Focus (examples)	Allowed actions	Model / params
UI suggestion	increase font sizes; enable contrast mode; stronger button borders; tooltips	<code>increase_font_size</code> , <code>increase_contrast</code> , <code>increase_button_border</code> , <code>show_tooltip</code>	<code>gemini-2.5-flash-lite</code> , temp=0.2, timeout=15s, thinking.budget=0
Geometry	spacing; button size; slider size; simplify layout	<code>increase_button_size</code> , <code>increase_slider_size</code> , <code>adjust_spacing</code> , <code>simplify_layout</code>	<code>gemini-2.5-flash-lite</code> , temp=0.2, timeout=15s, thinking.budget=0
Input	mode switching; trigger buttons	<code>switch_mode</code> , <code>trigger_button</code>	<code>gemini-2.5-flash-lite</code> , temp=0.2, timeout=15s, thinking.budget=0
Validator	consolidate + validate all suggestions	<i>all actions above</i>	<code>gemini-2.5-flash</code> , temp=0.3, timeout=30s, thinking.budget=-1

Table 5.3: MA-SIF agents, focus, allowed actions, and model settings (runtime configurable).

Runtime defaults:

- **Default model settings:**
 - UI/Geometry/Input use `gemini-2.5-flash-lite` with a temperature of 0.2, a 15s timeout and thinking budget of 0 (no thinking).
 - Validator uses `gemini-2.5-flash` with a temperature of 0.3, a 30s timeout and thinking budget of -1 (dynamic).
- **Whitelisting:** Each agent is restricted to an *allowed action set*; Validator allows the union set only.
- **Schema enforcement:** Responses must match the JSON schema (`response_json_schema`); non-conforming replies trigger fallback.
- **Low temperature + narrow prompts:** Specialist prompts reduce creative drift and injection surface.
- **Timeouts:** Per-agent time budgets ensure the loop returns within interaction thresholds.

5.5.4 LLM invocation (Gemini)

For each event, the backend invokes Gemini per specialist agent (UI, Geometry, Input) and then once for the Validator, in sequence. Each call is single-shot (no streaming) with a strict JSON response contract. Requests are composed as follows:

Request composition:

- **Context blocks:** `event_json`, `profile_json` (accessibility flags, preferences), and `history_json` (recent 10 interactions). These are injected into the Agent’s prompt at runtime, as mentioned earlier.
- **Agent preamble:** a short, role-specific instruction and focus items plus an explicit *allowed action set* (whitelisting).
- **Output contract:** instructed to return only a JSON envelope `{"adaptations": [...]}` (no prose).

Generation configuration (per call).

- **Model/temperature:** UI, Geometry, Input use `gemin-2.5-flash-lite` at $T=0.2$; the Validator uses `gemin-2.5-flash` at $T=0.3$ (see `sif_config.json`).
- **Response format:** `response_mime_type=application/json` with a **JSON Schema** that enforces types, fields, and that either `value` (numeric) *or* `mode` (categorical) is present.
- **Budgets:** per-agent timeout ≈ 15 s; Validator timeout ≈ 30 s. Calls are sequential by design to keep rate-limits predictable and simplify partial results.
- **Determinism:** low temperatures, single candidate, narrow prompts; agents propose domain-scoped adaptations, Validator consolidates/filters to the final set.

Failure handling. If a specialist agent times out or returns invalid JSON, its suggestions are omitted; the Validator runs on whatever is available. If validation still fails schema checks or time outs, the backend emits the available adaptation or conservative rule-based adaptations as a fallback.

5.5.5 Structured outputs and guardrails

To reduce hallucinations and schema drift, the backend requests JSON-typed responses from the LLM with an explicit schema and allowed actions, as described earlier in Chapter 4. The prompt defines required fields, expected types, and a one-of constraint that demands either a numeric `value` or a categorical `mode`. Agent prompts are intentionally narrow, which improves determinism. The validator prompt is broader, since it must adjust conflicting suggestions and justify final choices. Despite these controls, invalid outputs still occur occasionally. However, the schema can be enforced at two distinct levels: by the agents and by the validator. This dual-level approach makes it easier for the validator to ensure that the agents’ adaptations are compliant. Furthermore, it has a defensive layer that falls back to simple rules if all LLM validation fails. The output schema will produce an adaptation of the following format:

```
{"adaptations": [{"action": "...", "target": "...", "value\mode": "...", "reason": "...", "intent": "..."}]}
```

5.5.6 Rule-based fallback and resilience

A lightweight rule engine acts as a safety net when LLM calls time out or the provider is unavailable. This rule engine comes into action when all the suggestion agents fail (no output). It covers essential accessibility behaviors, for example increasing button size after a miss-tap, switching to voice for users flagged as motor-impaired, or enabling high-contrast mode for visually impaired profiles. These rules are intentionally conservative, they guarantee progress without surprising the user, and they keep the implementation usable in environments with unstable connectivity.

Fallback rules (summary)

Table 5.4: Rule triggers and conservative adaptations (fallback path).

Trigger	Adaptation(s)
Miss tap (<code>event_type="miss_tap"</code>) or slider miss	<code>increase.button.size</code> on target (default 1.3); intent=“difficulty hitting target”.
Profile indicates motor impairment	<code>increase.button.size</code> on all; intent=“Improve accessibility”.
Profile indicates visual impairment	<code>increase.contrast</code> on all; intent=“improve visibility”.
Voice command detected (<code>event_type="voice"</code>)	<code>switch_mode</code> to voice on all.

5.5.7 Heatmap Analysis

Due to time constraints, heatmap analysis is not implemented. This is a potential area for future work, as understanding user interaction patterns and frequent touch points, could further enhance the system’s adaptability. Heatmap analysis has a big potential in the repositioning of UI elements based on the

heatmap data, this could improve accessibility for users with specific needs like gesture-based inputs. Since the repositioning of elements is not supported in this first iteration, the focus remains on immediate adaptation strategies, which made heatmap analysis out of scope.

5.5.8 Latency, partial results, and error handling

The WebSocket loop is designed to return something useful as quickly as possible. Agent calls run in sequence within short time budgets. If one agent fails to respond, the validator operates on the remaining suggestions rather than waiting indefinitely. The backend aims to keep per-event processing below the threshold where users notice a lag on interaction, which is important for accessibility, particularly when enlarging targets immediately after an error. In practice, most adaptations are returned quickly by the smaller suggestion agents, while validation can become the slowest step in complex scenes. When validation exceeds its budget the backend returns the best available subset, then continues to append the event to the user's history so future interactions benefit from the context.

5.5.9 Security and CORS considerations

During development the backend enables permissive CORS to simplify local testing across platforms. Profiles are keyed by `user_id` rather than personal data. For production deployment, stricter origins, authentication, and encryption would be required. These measures are outside the scope of this prototype, but the separation of concerns in the current design makes them straightforward to add.

5.5.10 Summary

In its current form the backend delivers a complete adaptation pipeline: events arrive over WebSocket, profiles and short history windows are loaded from MongoDB, MA-SIF produces structured suggestions, a validator consolidates them, and the result is returned to the frontend within a single interaction loop. When LLM reasoning is unavailable, conservative rules ensure the interface remains usable. This combination of multi-agent reasoning, strict schemas, and rule-based fallbacks gives the system both flexibility and reliability, which is essential for accessibility-focused adaptations.

5.6 User Profile and Context Implementation

The user profile and context subsystem was implemented as a dedicated data service in the SIF backend, designed to persist accessibility needs, interaction preferences, baseline UI configurations, and a capped history of recent events. MongoDB serves as the primary storage layer, with the `profiles` collection indexed on `user_id` as described earlier for constant-time retrieval during event processing.

When a new event is received via the WebSocket (`/ws/adapt`), the backend queries the profile store using the supplied `user_id`. If no profile exists, the profile will be created and added to the MongoDB using `insert_one` before continuing, otherwise the profile will be retrieved using `find_one` or updated using `update_one`, ensuring that all adaptation decisions are made in a contextualized environment. This retrieval step is synchronous, guaranteeing that the most recent committed profile is available to the reasoning pipeline before any LLM or rule-based evaluation occurs. Interaction history is maintained using MongoDB's `$push` with `$slice` operators to append the incoming event while capping the array length at 10 entries for efficiency. This rolling history provides the SIF agents with temporal context, enabling progressive personalization; for example, recognising a pattern of repeated miss-taps and proactively switching to voice mode. Updates to the history are performed asynchronously to avoid blocking real-time adaptation.

Profile documents are structured as JSON as described in Chapter 4, containing four key sections: `accessibility_needs` (boolean capability flags), `input_preferences` (preferred modality and fallback order), `ui_preferences` (default font size, button scale and more), and `interaction_history` (recent event log). This schema strikes a balance between simplicity and extensibility, allowing new fields to be added without migration overhead.

To optimise performance and safety, all profile mutations are atomic, relying on MongoDB transactions to prevent race conditions when simultaneous events and updates occur. Adaptation logs are stored separately in the `logs` collection and mirrored to a local `adaptation_log.jsonl` file, supporting offline analysis and reproducibility of evaluation results. Furthermore, if a profile update is in-flight during an

event, the backend uses the latest committed profile, mitigated by client-side checks (waiting for `POST /profile` success) and server-side transactions.

This implementation ensures that every adaptation decision, whether produced by a static rule or the multi-agent LLM pipeline, is grounded in the user’s persisted profile and immediate interaction context, enabling consistent, personalised, and stateful UI behavior across sessions.

5.7 Dynamic Adaptation Mechanisms Implementation

The dynamic adaptation mechanism is the final stage in the adaptation pipeline, where decisions made by the backend are translated into immediate and visible changes in the user interface. In the current implementation, this process is tightly integrated with Flutter’s reactive widget system, allowing adaptation actions to be applied without forcing full UI rebuilds or navigation resets.

When the backend sends an adaptation list over the WebSocket connection, the frontend parses each action and routes it to the relevant UI element. Actions are defined in the strict JSON schema described earlier, containing the `action` type, a `target` identifier, a `value` or `mode`, and a human-readable `reason` and `intent` it inferred from the adaptation. This standardisation allows the same adaptation handler to process diverse actions without requiring modality-specific logic.

The framework applies a predefined set of accessibility-oriented actions (see Table 5.1 for the concrete widget mappings used by the Flutter implementation). These adaptation actions form the cornerstone of the framework’s accessibility-driven approach, addressing diverse user needs in the Adaptive Smart Home Controller. Each action is carefully designed to align with WCAG 2.1 guidelines, ensuring inclusivity for motor-impaired users (e.g., larger buttons/sliders, increased spacing), visually impaired users (e.g., high-contrast modes, highlighted borders), and hands-free users (e.g., tooltips, mode switching).

5.7.1 Application Mechanics (State, Animation, Ordering)

- **State deltas.** Element-scoped actions update per-element maps (e.g., `buttonScales[target]`, `elementBorders[target]`), while global actions flip flags or theme variables (`increase_contrast`, `switch_mode`).
- **Animation semantics.** Size and spacing changes use lightweight implicit animations (`AnimatedScale`, `AnimatedContainer`) with short durations to make changes perceivable without disrupting interaction.
- **Ordering.** Adaptations are applied in the order received; the backend validator guarantees a coherent, non-overlapping set. On receipt, the active card’s glow border stops and controls are re-enabled.

The application of adaptations begins with a lookup to determine whether the `target` element exists in the current view. If it does, the corresponding widget state is updated directly. For example, an `increase_button_size` action adjusts the scale factor property of the button widget, and an `increase_font_size` action updates the text style parameter. Where appropriate, changes are animated using Flutter’s `AnimatedScale` or `AnimatedContainer` to make the transition noticeable without distracting the user. This animation step is particularly important for accessibility, as it helps the user understand that the interface has been modified intentionally. The loading indicator displayed around the card being interacted is now also triggered to stop playing.

Adaptations that affect the entire interface, such as `increase_contrast` or `switch_mode`, are handled at the application theme level. Contrast adjustments update the color palette by replacing the primary and background colors with higher-contrast alternatives, while mode switches alter the active input modality, for example switching from touch to voice. These global changes are propagated across all widgets automatically through Flutter’s state management, ensuring consistency without manually updating each element (see table 5.5 for an overview).

Table 5.5: Scope of adaptations in the Flutter implementation.

Global (theme/state)	increase_contrast, switch_mode, simplify_layout, increase_font_size (all)
Element-scoped	increase_button_size, increase_button_border, increase_slider_size, adjust_spacing, show_tooltip, trigger_button

Not all adaptations are implemented with live modality inputs. For demonstration purposes, actions triggered by voice or gesture events are generated from simulated events in the frontend’s test panel in each device card, like was mentioned earlier. However, these simulated actions follow the same processing path as real events, which means that integrating actual input sources in the future will require no changes to the adaptation mechanism itself.

5.7.2 Conflicts and Unknown Actions

The validator resolves conflicts server-side; the frontend applies the resulting set atomically in a single frame. If duplicate actions target the same element, the last write wins within the frame. Unknown or ill-typed actions are logged and ignored, preserving UI stability.

5.7.3 Real-Time Adaptation Example

To illustrate the framework’s adaptation capabilities, consider a scenario where a visually impaired user with voice preference attempts to interact with the Smart Home Controller. The user profile contains the following settings:

```

1 {
2   ...
3   "accessibility_needs": { "visual_impaired": true },
4   "input_preferences": { "preferred_modality": "voice" },
5   "ui_preferences": { "font_size": 16, "button_scale": 1.1 }
6 }
```

Listing 5.7: User profile for visually impaired user, shortened

Event Trigger: The user attempts to activate the lamp by tapping its button but misses the target due to the button’s insufficient size relative to their visual and motor coordination needs. This generates a `miss_tap` event on the `lamp` target, which is captured by the frontend and sent through the Input Adapter Layer to the SIF backend.

SIF Reasoning Process: The backend combines the miss-tap event with the user’s profile and recent interaction history. The MA-SIF agents analyze the situation:

- **UI Agent:** Recognizes visual impairment flags and suggests contrast and font improvements
- **Geometry Agent:** Identifies the miss-tap pattern and recommends button enlargement and border enhancement
- **Input Agent:** Notes voice preference and triggers the intended action while suggesting mode awareness
- **Validator Agent:** Consolidates suggestions and ensures compatibility

Adaptation Response: The backend returns five coordinated adaptations:

1. **Font size increase (lamp):** Scale text by 1.2× for better readability given visual impairment
2. **Button size increase (lamp):** Enlarge the lamp button by 1.2× to reduce future miss-taps

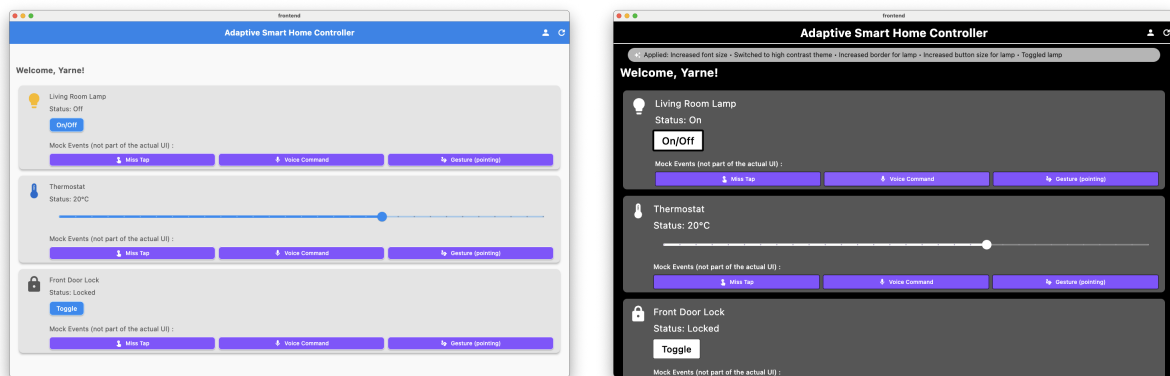
3. **Contrast enhancement (global):** Switch to high-contrast mode across the entire interface
4. **Button border enhancement (lamp):** Strengthen the lamp button border for better visual definition
5. **Direct button activation (lamp):** Trigger the lamp immediately via voice command, fulfilling the user's intent despite the miss-tap

Validator Agent Reasoning: The Validator accepted all (5 out of 5) adaptations from the suggestion agents, following this reasoning:

- **Font size increase (lamp):** User has visual impairment and has previously missed taps on the lamp button, suggesting difficulty in seeing the target.
- **Contrast enhancement (global):** User has visual impairment, and increasing contrast can improve visibility of UI elements.
- **Button border enhancement (lamp):** User has visual impairment and has previously missed taps on the lamp button, suggesting difficulty in accurately targeting the button.
- **Button size increase (lamp):** The user has a visual impairment and has recently missed tapping the lamp button twice. Increasing the button size will make it easier to interact with.
- **Direct button activation (lamp):** User has repeatedly missed tapping the lamp button, indicating a potential difficulty with precise touch interactions. The current miss-tap command to turn on the lamp suggests a preference for hands-free operation. Triggering the button via voice command is a suitable adaptation.

Visual Impact: The transformation (see figure 5.8) demonstrates the framework's accessibility-focused approach. The originally small lamp button (standard size) becomes significantly more prominent through size scaling and border enhancement. Text elements throughout the interface are enlarged to improve readability, while the high-contrast theme ensures better visibility for users with visual impairments. The lamp is also activated immediately, fulfilling the user's original intent despite the initial miss-tap.

Personalization Benefits: This example illustrates how the framework moves beyond generic accessibility settings to deliver contextual, profile-aware adaptations. Rather than applying a one-size-fits-all solution, the system combines the user's declared visual impairment with their interaction behavior (miss-tap) and stated preferences (voice modality) to produce a coordinated set of improvements that address both immediate needs (completing the lamp activation) and future interactions (persistent size and contrast improvements). The adaptation occurred within the framework's typical latency range (~13 seconds median), allowing the user to continue interacting while improvements are applied progressively. This demonstrates the practical viability of near real-time, AI-driven adaptation for accessibility enhancement.



Before: standard interface with small targets

After: enlarged targets, enhanced contrast, stronger borders

Figure 5.8: Real-time adaptation applied after miss-tap event for visually impaired user.

5.8 Backend Injection Interface

5.8.1 Purpose and Scope

This lightweight Flutter tool exposes a direct interface to the SIF backend for debugging and evaluation. It lets a researcher inject or edit *profiles* and *events*, send them to the backend over WebSocket/HTTP, and immediately inspect the returned `{adaptations:[...]}` along with the server-side interaction history per user. It is separate from the end-user UI and optimized for fast iteration.

5.8.2 Architecture and Data Flow

On launch the app opens a WebSocket to `ws://localhost:8000/ws/adapt` and listens for JSON replies containing an `adaptations` array. Profile updates are sent via HTTP POST `/profile`, while the history panel pulls entries from GET `/full_history`. Incoming WS frames are decoded and the `adaptations` list is rendered; after each reply the tool refreshes history via HTTP to keep the view in sync.

5.8.3 Controls: Profiles and Events

The left column edits the *profile* (free-form JSON or prefilled “Motor Impaired” / “Hands-Free”), visualizing key blocks like `accessibility_needs` and `input_preferences`. The right column edits the *event* (free-form JSON or predefined “Miss-Tap on Play”, “Voice Play Command”, “Gesture Point”). Both editors pretty-print JSON and show an iconized summary (user, modality, target). This enables fast A/B of context vs. input without leaving the app (see figure 5.9 for an example).

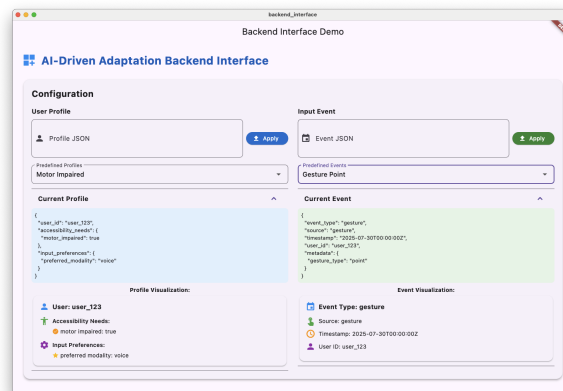


Figure 5.9: Backend Injection Interface: configuration and live JSON visualization.

5.8.4 Adaptation Response View

Pressing *Get Suggestions* sends the current event (after posting the profile), shows a spinner, then lists each adaptation with an action-specific icon and a compact, human-readable summary (e.g., “Switch to voice mode” or “Increase contrast”). The mapping covers common actions such as `switch_mode`, `increase_contrast`, `trigger_button`, `simplify_layout`, and geometry actions. Unknown actions fall back to a generic renderer, keeping the tool resilient to backend experiments (see figure 5.10 for an example).

5.8.5 Interaction History Panel

The bottom card fetches `history` from the backend and renders each entry (timestamp, user, event type/source) plus the associated adaptations. This provides a quick audit trail to verify that profile changes affect subsequent suggestions and to spot repeated error patterns (e.g., consecutive miss-taps). A “Clear History” button resets the local view; history retrieval is refreshed after each WS reply to keep causality visible.

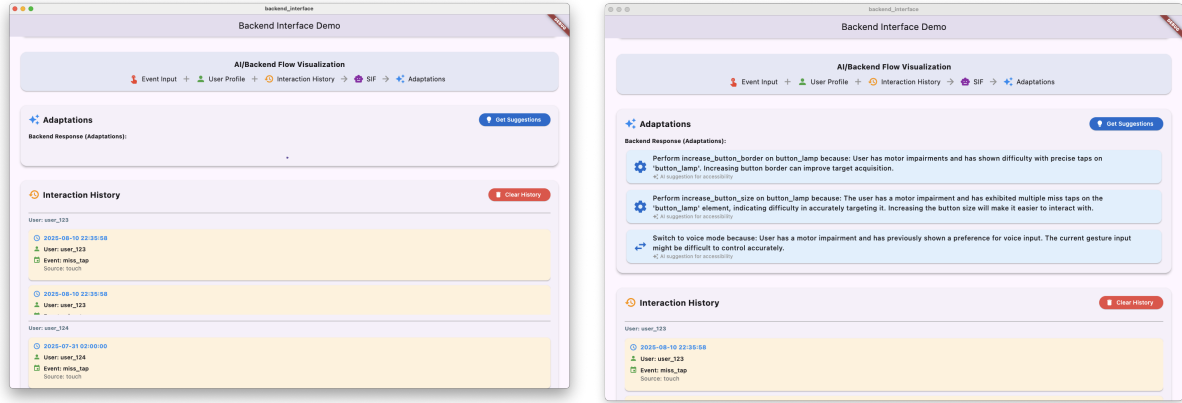


Figure 5.10: Example adaptation list before/after a pointing gesture event.

5.8.6 Operational Notes and Limitations

The app assumes a running backend at `localhost:8000`. Profile JSON is posted verbatim; there is no schema validation beyond backend checks. The WS channel is established once on init; a commented re-connection helper is included for future hardening. Timeouts and transport errors are surfaced minimally to keep the UI uncluttered. This is created as a developer tool and is not intended for end-users.

5.9 Cross-Platform SwiftUI Example

5.9.1 Purpose and Scope

This minimal SwiftUI app demonstrates that a thin, platform-native adapter plus a small view layer are sufficient to integrate with the SIF backend. It mirrors the Flutter contract: create (or fetch) a user profile over HTTP, stream events over WebSocket, and apply the returned `{adaptations:[...]}` envelope to the UI state. The example targets iOS/macOS (Swift 5.7+, iOS 15+) and may require ATS exceptions for local HTTP/WS during development.

5.9.2 Adapter and Transport

The `AdaptiveUIAdapter` is an `ObservableObject` that (i) ensures the profile exists via `GET /profile/{user_id}` and `POST /profile`, (ii) opens a `URLSessionWebSocketTask` to `/ws/adapt`, and (iii) exposes `sendEvent()` and an `onAdaptations` callback for the UI. Incoming frames are decoded from a wrapped `Envelope{adaptations:[...]}` into an array of `Adaptation`. On receive errors it reconnects after a short delay, keeping the session resilient for local testing and debugging. Events are encoded with `JSONEncoder.convertToSnakeCase`, matching the backend’s `snake_case` contract.

5.9.3 Minimal UI and State Mapping

`AdaptiveUIApp` wires the adapter into the environment and renders a single `ContentView` that shows a lamp card, a few “send event” buttons, and a live list of the most recent adaptations (see figure 5.11 for a before-after example). The view holds a small set of reactive states `buttonScale`, `borderWidth`, `highContrast`, `fontScale`, `spacing`, and `tooltip`, that the `apply()` function updates based on `action/target/value/mode`. Concretely: `increase.button_size` multiplies `buttonScale`, `increase_font_size` scales text, `increase_contrast` toggles a high-contrast color scheme, `increase.button_border` sets a visible outline, `adjust_spacing` scales inter-control spacing (clamped), `show_tooltip` surfaces the model’s reason string, and `trigger_button` toggles the lamp state. This mirrors the Flutter mapping but proves the contract is UI-framework agnostic.

5.9.4 Event Injection

Three buttons generate representative events: a miss-tap (`event_type="miss_tap"` with coordinates), a simple voice command (`event_type="voice"` with `metadata.command`), and a local tap. The miss-tap

and voice paths serialize to JSON and are sent over WS; the tap path is stubbed to flip local state, emphasizing that production UIs would simply send the tap event the same way.

5.9.5 Usage and Limitations

On launch, the adapter creates a minimal default profile when needed, connects WS, and begins listening. Pressing “Miss-tap” or “Voice” triggers a backend roundtrip; the adaptation list below the card updates immediately and `apply()` mutates view state. The demo intentionally omits persistence and complex navigation; its goal is to show that the SIF contract (profile endpoints, WS envelope, action names) can be integrated cleanly into SwiftUI with \approx a hundred lines of adapter+view code.

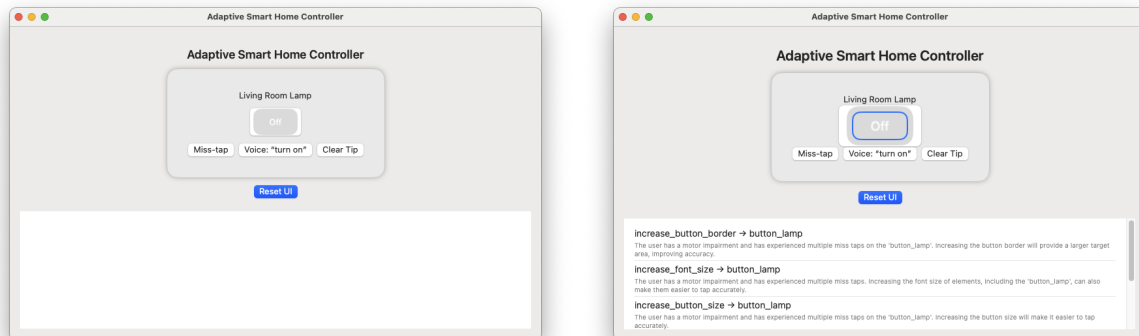


Figure 5.11: Example SwiftUI app before/after a miss_tap event.

5.10 Design Decisions

The design of the Adaptive Multimodal GUI Framework using LLMs reflects a series of deliberate choices aimed at balancing accessibility, performance, scalability, extensibility, and ease of integration. These decisions were made with the primary goal of delivering near real-time, personalised adaptations for motor-impaired, visually impaired, and hands-free users, while ensuring that the framework remains modular and adaptable to future platforms and domains.

5.10.1 Modularity Over Monolithic Design

Decision: The framework adopts a modular three-layer architecture (Frontend, Input Adapter, SIF Backend) with clear separation of concerns, connected through a standardized JSON contract.

Reasoning:

- **Flexibility:** Each layer can be updated or replaced independently, enabling deployment across different platforms such as Flutter, SwiftUI, or future platforms without reworking the entire system.
- **Extensibility:** The JSON contract in the Input Adapter Layer allows new modalities (e.g., eye tracking) to be integrated without modifying backend logic.
- **Developer accessibility:** Modularity simplifies integration, requiring minimal code for event handling and adaptation application.

5.10.2 WebSocket for Real-Time vs. HTTP for Batch Processing

Decision: The framework uses WebSocket (`/ws/adapt`) for real-time event processing and adaptation delivery, with HTTP (`/full_history`, `/profile`) for debugging, developer tooling and profile management.

Reasoning:

- **Low Latency:** WebSocket enables fast and bidirectional sending of data like adaptations (e.g., scaling a button after a miss-tap).

- **Reliability:** HTTP supports robust profile updates (`POST /profile`), ideal for non-real-time scenarios or debugging.
- **Accessibility:** Real-time feedback enhances usability for motor-impaired or hands-free users, where delays could disrupt interaction.

5.10.3 MongoDB for Persistent Storage

Decision: MongoDB is used for storing user profiles, interaction history, and adaptation logs, with `user_id` indexing and capped history arrays (10 events).

Reasoning:

- **Scalability:** MongoDB’s NoSQL design and indexing ensure fast queries for large user bases, critical for real-world deployment.
- **Flexibility:** JSON-like documents align with the framework’s JSON contract, simplifying profile/history storage.
- **Continuous Learning:** Retaining a limited history supports adaptive behavior, such as making permanent size adjustments after repeated miss-taps.

5.10.4 Rule-Based Fallback with LLM Reasoning

Decision: SIF combines rule-based logic (e.g., `if miss_tap then increase_size`) with LLM reasoning for creative adaptations, with rules as a fallback for LLM failures or time-outs, since LLMs can have high respond latencies.

Reasoning:

- **Reliability:** Rules ensure deterministic adaptations (e.g., button enlargement for miss-taps) when LLM responses are unavailable or hallucinate.
- **Novelty:** LLM enables context-aware, proactive suggestions (e.g., `switch_mode: voice` for hands-free users), advancing beyond static rules.
- **Accessibility:** Hybrid approach ensures consistent support for motor-impaired, visually impaired users (e.g., high-contrast text).

5.10.5 Multi-agent LLM reasoning (MA-SIF) vs single-agent LLM reasoning (normal SIF)

A key architectural decision in the framework is the adoption of multi-agent LLM reasoning (MA-SIF) over a single agent LLM approach (normal SIF). In the single agent SIF model, one LLM is responsible for interpreting all input events and generating adaptation actions. While this simplifies integration and reduces system complexity, it can limit the granularity and specialization of adaptation logic, especially as the diversity of user needs and input modalities grows.

MA-SIF, by contrast, distributes reasoning across multiple specialized LLM agents, each focused on a distinct domain such as UI adaptations, geometry/layout changes, input modality management, and validation. This separation of concerns enables each agent to leverage tailored prompts, domain-specific knowledge, and focused reasoning strategies, resulting in more nuanced and context-aware adaptation suggestions. The validator agent further ensures that outputs from other agents are coherent, non-conflicting, and accessibility-compliant.

The multi-agent approach offers several advantages:

- **Scalability:** New agents can be added to address emerging modalities or adaptation domains without disrupting existing logic.
- **Extensibility:** Prompts and allowed actions can be updated independently for each agent, supporting rapid iteration and domain-specific improvements.
- **Robustness:** Specialized agents reduce the risk of LLM hallucinations or conflicting adaptations, as validation is enforced before application.

- **Personalization:** Agents can incorporate user profiles and history more effectively, enabling targeted adaptations for motor-impaired, visually impaired, or hands-free users.

5.11 Implementation Challenges and Solutions

Developing the Adaptive Smart Home Controller exposed several practical challenges, ranging from LLM-specific issues to general real-time system concerns. These were addressed through a mix of architectural decisions, fallback mechanisms, and compromises designed to keep the prototype functional and reliable under varying conditions.

5.11.1 LLM reliability and output consistency

One of the most persistent challenges was ensuring that the multi-agent LLM pipeline returned valid, schema-compliant output. Despite strict prompts and an explicit allowed-actions list, hallucinations still occurred, producing unsupported actions, targeting non-existent elements, or returning unreasonable values such as excessively large scaling factors. These errors risked breaking the layout or producing jarring visual changes. To mitigate this as described earlier, a validator agent was introduced to consolidate and correct suggestions before they reached the frontend. However, this validation step added extra latency and still could not guarantee complete protection against invalid values, making additional frontend-side checks a sensible next step for future versions.

5.11.2 Performance under real-time constraints

Adaptation latency was a critical factor for usability, particularly in accessibility scenarios where delayed feedback can reduce trust in the system. Smaller LLM models were assigned to the suggestion agents to keep their execution time within fractions of a second, while the validator, which required more context and reasoning, used a larger model with a higher timeout budget. Even so, occasional slowdowns occurred due to network latency or provider-side delays. The backend was therefore designed to apply partial results if all agents did not respond in time, ensuring that at least some adaptations reached the user quickly.

5.11.3 Safeguards against malicious or replay attacks

In its current form, the system does not include strong safeguards against replay attacks or intentionally crafted events designed to trigger disruptive adaptations. This is acceptable in a controlled research environment but would require strict validation and authentication for production deployment. Measures such as signing WebSocket messages, verifying sequence numbers, and rejecting stale or malformed events would be necessary to prevent exploitation as well as extra protections in the adapter or frontend. Furthermore is `user_id` hijacking, where an attacker submits events under a different user's profile possible, since no `user_id` checks are done on returning adaptations. Future mitigations such as implementing `user_id` verification and event signing will be necessary to address these vulnerabilities. This was intentional for the scope and development time constraints of this research.

5.11.4 Testing with incomplete modalities

Live gesture tracking and voice inputs were not integrated in this iteration, which meant that related adaptations had to be tested using simulated events. While this allowed the adaptation logic to be validated end-to-end, it did not account for the noise, recognition errors, or latency introduced by real input devices. Future work will need to focus on replacing these simulated events with actual input sources to fully evaluate performance in realistic conditions.

5.11.5 Security and trust boundaries

In its current implementation state, the backend accepts connections from any origin and does not require authentication for profile creation or event submission. Currently a pure localhost setup is implemented for the entire backend. This was a deliberate decision to speed up testing across platforms, but would need to be replaced with stricter CORS rules, token-based authentication, and access control in production. Without these measures, the system is vulnerable to unsolicited adaptation requests from external sources.

Furthermore, some config files such as `sif_config.json` do not have the necessary validation checks in the backend, while the validator agent is necessary, nothing stops it from being removed.

5.12 Chapter Summary

This chapter has detailed the implementation of the Adaptive Smart Home Controller as the primary proof-of-concept for the multimodal AI-driven GUI framework. The system was realised as a three-layer architecture, consisting of a Flutter-based frontend for rendering and applying adaptations, an Input Adapter Layer for standardising and transmitting events, and a FastAPI-based backend implementing multi-agent Smart Intent Fusion with MongoDB for persistent profile and history storage.

Key implementation aspects included the construction of a dynamic adaptation pipeline capable of applying size, contrast, modality, and content changes in real-time, and the use of a strict JSON schema to maintain consistency between layers. While some modalities such as voice and gesture inputs were simulated for demonstration purposes, the system was designed so that integrating real devices will require minimal architectural changes.

The prototype also faced practical challenges, including LLM output validation, latency management, security considerations, and the absence of certain safeguards against malicious or replayed events. These were addressed through a combination of a validator agent, conservative rule-based fallbacks, partial-result handling, and a modular design that isolates critical components.

In its current state, the implementation demonstrates that the framework can deliver personalised, accessibility-focused adaptations in a responsive and modular manner, even under the constraints of a prototype environment. The next chapter evaluates this implementation through a feasibility study, assessing its responsiveness, adaptability, and practical potential in simulated real-world scenarios.

Chapter 6

Feasibility Study

6.1 Study Overview

The evaluation follows a feasibility study design to check whether the multimodal AI-driven adaptation pipeline works end to end under realistic interaction traces, and whether it does so in a way that is useful for accessibility. Concretely:

1. **Adaptation validity:** Does the backend return schema-valid adaptation objects across different profiles and modalities?
2. **Accessibility relevance:** Do adaptations match user needs (motor, visual, hands free) and observed errors (for example `miss_tap`, `slider_miss`)?
3. **Latency:** What is end-to-end response time (p50/p90/max) under the chosen backend configuration?

6.2 Methodology

Profiles and runs. Six user profiles (P0–P5) were selected to cover baseline, single-need, and combined needs:

- **P0 (Baseline):** No specific accessibility needs.
- **P1 (Motor):** Larger touch targets; keyboard preferred.
- **P2 (Visual):** Larger font and higher contrast preferred.
- **P3 (Hands free):** Voice and gesture preferred.
- **P4 (Motor + Hands free):** Combined motor support and voice preference.
- **P5 (Visual + Motor):** Combined visual and motor support.

For each profile, two runs were executed to allow a small history to build. Each run used the same seven-event sequence (14 events per profile, 84 total); see Table 6.1.

Table 6.1: User profile mapping for evaluation runs.

Profile	Declared needs
P0	Baseline (no declared needs)
P1	Motor, keyboard preferred
P2	Visual, larger font
P3	Hands free preferred, voice modality
P4	Motor + Hands free, voice modality
P5	Visual + Motor, voice modality

Event suite. The suite mixes errors and successful attempts across modalities and targets (lamp, thermostat, lock):

- `miss_tap` on buttons (lamp, lock),
- `slider_miss` with overshoot (thermostat),
- voice commands (turn on, unlock, adjust),
- `gesture` inputs (point/select).

This combination triggers typical accessibility challenges (tap precision, slider control, and mode switching).

Backend configuration and metrics. All traces used the MA-SIF (balanced) configuration (Table 6.2). Each event’s response was checked for schema validity; the classification path and client-side latency (p50/p90/max) were logged. Analysis scripts flatten JSONL logs to CSV and compute objective metrics.

Table 6.2: MA-SIF (balanced) agent configuration.

Agent	Model	Thinking budget	Temperature	Timeout (s)
UI	gemini-flash-lite	0	0.2	15
Geometry	gemini-flash-lite	0	0.2	15
Input	gemini-flash-lite	0	0.2	15
Validator	gemini-flash	-1	0.3	30

Hardware and Environment: All experiments in this thesis were executed on a single development machine in a controlled home environment (none to low network contention). Latency is end-to-end from event send to response received on the client. Hardware configuration:

- **Hardware:** MacBook Pro, M1 Max 10-core CPU, 32 GB RAM
- **Software:** macOS 15.6 (Sequoia), Python 3.9+
- **Environment:** Localhost + Gemini external API calls

Operational definitions. **Schema-valid** means the response conforms to the adaptation schema (valid action names, parameters in bounds, resolvable UI targets, and consistency checked by the Validator). **Latency** is measured as $\Delta t = t_{\text{recv}} - t_{\text{send}}$.

Objective scoring rules. All quality metrics are computed deterministically from the logs using fixed contracts and allowed actions. Two lookup tables are used:

(i) **Action** \rightarrow **Accessibility Category**

Table 6.3: Mapping of adaptation actions to accessibility categories.

Action	Accessibility Category
<code>increase.button.size</code>	Motor
<code>increase.button.border</code>	Motor
<code>increase.slider.size</code>	Motor
<code>adjust.spacing</code>	Motor
<code>increase.font.size</code>	Visual
<code>increase.contrast</code>	Visual
<code>switch.mode</code> (voice or gesture)	Hands-free
<code>trigger.button</code> (from voice intent)	Hands-free

(ii) Error Event \rightarrow Acceptable Corrective Adaptations**Table 6.4:** Mapping of error events to acceptable corrective adaptations.

Error Event	Acceptable Corrective Adaptations
miss_tap	increase_button_size, increase_button_border, adjust_spacing, switch_mode:voice
slider_miss	increase_slider_size, adjust_spacing
voice (command)	switch_mode:voice, trigger_button (if intent clear)
gesture (point/select)	switch_mode:gesture, trigger_button (if intent clear)

These tables are grounded in the schema and rule-based fallbacks, which makes the assessment replicable and objective.

6.3 Accessibility and Design Quality

This section reports effectiveness before latency, using objective, log-derived metrics. Effectiveness is assessed by quantifying how well the adaptation pipeline produces accessibility-relevant, profile-aligned, and internally coherent UI changes in response to realistic interaction traces. Metrics such as Profile–Action Alignment (PAA), Error to Response Appropriateness (ERA), Mode Enablement for Hands free (MEH), and Design Coherence Index (DCI) are computed directly from adaptation logs, enabling reproducible and objective evaluation. These results provide a detailed view of the system’s accessibility impact and personalization quality prior to analyzing performance characteristics such as latency.

6.3.1 Profile–Action Alignment (PAA)

Definition. PAA is the share of accessibility actions whose category matches the active profile’s needs (motor, visual, hands free):

$$\text{PAA} = \frac{\#\{\text{cat}(a) \in \text{needs}(p)\}}{\#\{\text{accessibility actions}\}} \times 100\%.$$

Result. Across non baseline profiles (P1–P5), and using the top five actions per profile, PAA is approximately **55%**. Combined-need profiles show higher alignment (P4 80.00%, P5 77.42%) than single-need profiles that are underrepresented by the event suite (P2 34.69%, P3 28.00%); see Table 6.6.

6.3.2 Error to Response Appropriateness (ERA)

Definition. ERA is the share of error or modality events that received at least one acceptable corrective adaptation:

$$\text{ERA} = \frac{\#\{\text{error events with } \geq 1 \text{ corrective adaptation}\}}{\#\{\text{error events}\}} \times 100\%.$$

Result. ERA ranges from **50.00%** to **64.29%** per profile (for example P5 64.29%), indicating that about one half to two thirds of these events contain a corrective action from the predefined set (Table 6.6).

6.3.3 Mode Enablement for Hands free (MEH)

Definition. MEH is the share of events in which a hands free path is explicitly enabled or used (switch_mode:voice|gesture or voice trigger_button).

$$\text{MEH} = \frac{\#\{\text{events with switch_mode:voice|gesture or accepted trigger_button via voice}\}}{\#\{\text{events}\}} \times 100\%.$$

Result. MEH equals **57.14%** for hands free profiles P3 and P4, indicating frequent enablement of non touch paths; values are lower for profiles where hands free is not a primary need (Table 6.6).

6.3.4 Design Coherence Index (DCI)

Definition. DCI quantifies the absence of contradictions and duplicates within a single response:

$$\text{DCI} = 1 - \frac{\#(\text{conflicts} + \text{duplicates})}{\#\text{suggestions}}.$$

Result. DCI is high across profiles (mean ≈ 0.995), which indicates internally consistent, non redundant suggestions (Table 6.6).

6.3.5 Global accessibility share and WCAG coverage

Overall, **97.51%** of all actions are accessibility targeted (352/361), see Table 6.7. The policy level coverage in Table 6.8 shows that proposed adaptations address SC 2.5.5 (Target Size), SC 1.4.3 (Contrast), and SC 1.4.4 (Resize Text). This is coverage, not a conformance audit.

6.3.6 Motor Benefit Proxy (Fitts law, analytic)

Rationale: When targets are enlarged by a factor s (for example $s=1.5$), the Fitts law index of difficulty changes from $\log_2(\frac{D}{W} + 1)$ to $\log_2(\frac{D}{sW} + 1)$. For representative D/W ratios, a $1.5\times$ size increase yields the following theoretical reductions:

Table 6.5: Fitts law proxy: relative reduction in index of difficulty for $s=1.5$.

D/W	2	4	8	12	16
Reduction (%)	22.9	19.3	16.0	14.3	13.3

This gives an interpretable proxy for expected motor effort reduction when applying `increase_button_size` or `increase_slider_size`.

Table 6.6: Objective accessibility and design quality metrics by profile.

Profile	PAA (%)	ERA (%)	MEH (%)	DCI	Acc. actions
P0	0.00	50.00	57.14	0.982	61/62
P1	50.88	57.14	57.14	0.986	58/60
P2	34.69	50.00	42.86	1.000	54/55
P3	28.00	50.00	57.14	1.000	57/58
P4	80.00	50.00	57.14	1.000	53/54
P5	77.42	64.29	57.14	1.000	69/72

Table 6.7: Overall accessibility targeted actions from the global action distribution.

Metric	Value
Total actions	361
Accessibility targeted	352
Share (%)	97.51
Aggregate aligned % across P1–P5 (top-5 actions)	63.72

Table 6.8: Policy level WCAG coverage addressed by proposed adaptations.

Profile	SC 2.5.5 Target Size	SC 1.4.3 Contrast	SC 1.4.4 Resize Text
P0	✓	✓	✓
P1	✓	✓	✓
P2	✓	✓	✓
P3	✓	✓	✓
P4	✓	✓	✓
P5	✓	✓	✓

6.3.7 Appropriateness by event type

This analysis reports *Error to Response Appropriateness* (ERA) per event type using the same acceptable corrective sets defined in this chapter. The 95% confidence intervals use the Wilson method.

Table 6.9: ERA by event type (higher is better).

Event	Count	ERA (%)	CI (binomial, 95%)	Notes
miss_tap	24	100.00	[86.20, 100.00]	motor fixes prevalent
slider_miss	12	100.00	[75.75, 100.00]	slider size/spacing
voice	36	97.22	[85.83, 99.51]	mode confirmation
gesture	12	8.33	[1.49, 35.39]	mode confirmation

Interpretation. ERA is very high for motor related errors: `miss_tap` and `slider_miss` both reach 100% with lower confidence bounds of 86.20% and 75.75% respectively. This indicates that the current rule and validation policy reliably proposes acceptable fixes such as larger targets, thicker borders, and increased spacing for these error modes. Voice commands also achieve a high ERA (97.22%), showing consistent mode confirmation and intent execution.

Gesture events show a low ERA (8.33%) with a wide interval due to the smaller sample size. This suggests that gesture handling is under specified relative to other modalities. In practice, gesture inputs often require additional disambiguation or confirmation to translate intent into a safe corrective action that passes validation. This requires more sophisticated gesture recognition and handling capabilities, which are outside the scope of this framework.

Overall, Table 6.9 shows that the policy handles motor errors and voice input effectively, while gesture interaction remains the main opportunity for improvement. This aligns with the personalization results in Table 6.6, which indicate lower alignment for hands free and visual oriented profiles when gesture is involved.

6.3.8 Stability across runs

To examine consistency, the adaptation sets of run 1 and run 2 are compared at the same event indices. Stability is measured as Jaccard similarity over action and target pairs at each index; values range from 0 (no overlap) to 1 (identical sets). Table 6.10 reports results for the first five events (E1 to E5) and the mean per profile.

Table 6.10: Stability across runs (Jaccard similarity of action and target sets).

Profile	E1	E2	E3	E4	E5	Mean
P0	0.11	0.12	0.22	0.20	0.14	0.16
P1	0.12	0.00	0.33	0.00	0.00	0.09
P2	0.11	0.00	0.40	0.00	0.14	0.13
P3	0.12	0.00	0.40	0.00	0.00	0.11
P4	0.40	0.14	0.40	0.17	0.17	0.26
P5	0.50	0.11	0.67	0.12	0.57	0.39

Interpretation. Stability is highest for the combined needs profile P5 (mean 0.39) and moderate for P4 (0.26). These values indicate that many core adaptations repeat between runs for users with multiple declared needs, which matches the global action mix (frequent motor and visual aids). Stability is lower for P1 to P3 (means 0.09 to 0.16). This reflects greater sensitivity to immediate context for single need profiles, where small differences in errors and interaction history can yield different yet still acceptable suggestions. The per event pattern also varies: E3 and E5 show higher overlap for several profiles, while E2 and E4 show lower overlap, which suggests that certain indices in the sequence encourage more exploratory adaptations.

Overall, the stability analysis suggests a balanced behaviour: adaptations remain consistent where combined needs call for recurring aids, and remain flexible where recent context should dominate. Very high stability is not always desirable because a responsive system should adapt to recent errors and history. The pattern in Table 6.10 is consistent with this goal: profiles with combined needs preserve a stable core of adaptations across runs, while single need profiles show more variation around that core.

6.3.9 Specificity via profile swap

A counterfactual analysis re-scores each response against the needs of a different profile (profile swap). Across swaps, PAA drops by **6.35** points on average (Table 6.11), which indicates that adaptations are tuned to the active profile rather than generic.

Table 6.11: Counterfactual PAA under profile swaps (mean over all swaps).

Profile	Original PAA (%)	Swap PAA (%)
P1	53.70	50.00
P2	34.69	52.55
P3	28.00	55.00
P4	80.00	43.00
P5	77.42	41.53

Interpretation. Specificity is strongest for combined needs (P4 and P5), where swap PAA falls by 37.00 and 35.89 points. This shows that the chosen adaptations align with those profiles and do not generalise to other needs. P1 shows a small drop (3.70 points), suggesting moderate specificity. P2 and P3 increase under swaps, which reflects the event suite bias toward motor and hands free aids: when their actions are re-scored against other profiles’ needs, more of them count as aligned. To raise specificity for single need profiles, the validator can apply stronger need weighting and add gentle penalties for off-need actions, and the event suite can include more visual and gesture specific signals so that on-need adaptations dominate.

6.3.10 Failure taxonomy (invalid cases)

By design, the SIF Validator returns only accepted adaptations; rejection reasons are not exposed in the response envelope. As a result, the logs do not support a fine-grained taxonomy of failure modes. Invalid responses are still measurable at the client by re-validating the final payload against the adaptation schema. Across all events, the invalid share is **15.48%** (derived from Table 6.15). Table 6.12 reports the per-profile distribution.

Table 6.12: Invalid responses by profile (derived from client-side schema check of final payloads).

Profile	Invalid (count)	Invalid (%)
P0	3	21.43
P1	2	14.29
P2	2	14.29
P3	4	28.57
P4	2	14.29
P5	0	0.00
Total	13	15.48

Interpretation. The highest invalid share appears for the hands free profile (P3), which aligns with other indicators that gesture and voice affordances are less fully specified in the current action menu and rules. The combined visual+motor profile (P5) shows no invalids, which is consistent with the strong alignment and stability observed for that profile.

6.3.11 Rule-only sufficiency (proxy)

For each event, the returned actions were checked against the acceptable corrective sets defined for that event type (Section 6). In **85.71%** of cases, a rule-only "fast path"¹ would have produced at least one acceptable action (aggregate from Table 6.9: $24+12+35+1 = 72$ hits over 84 events). This suggests that a staged design with a fast path followed by the Validator can lower median latency while preserving most coverage. A practical policy is to route obvious patterns (for example `miss.tap` → target enlargement and spacing; `slider.miss` → larger slider) through the fast path, while leaving ambiguous or gesture-driven cases to the Validator. This division keeps deterministic corrections quick and consistent, and reserves multi-agent reasoning for the smaller set of events where context and history matter most.

6.3.12 Rule baseline vs. LLM (minimal vs. maximal)

This analysis compares a deterministic rule baseline against the final LLM adaptations at the action-name set level. Two variants are used: (i) *minimal* rules that apply one canonical fix per event type, and (ii) *maximal* rules that include all acceptable fixes for that event type. Similarity is measured with Jaccard set overlap; the table also reports the baseline's own PAA. A top-5 action mask per profile matches the PAA style used elsewhere.

Table 6.13: Rule baseline vs. LLM by profile. Jaccard is set overlap of action names.

Profile	Jaccard (min)	Jaccard (max)	Exact (max)	Rule PAA (min %)	Rule PAA (max %)
P0	0.27	0.32	0.00	0.00	0.00
P1	0.27	0.37	0.00	42.86	55.56
P2	0.24	0.29	0.00	0.00	0.00
P3	0.26	0.34	0.00	57.14	44.44
P4	0.28	0.37	0.07	100.00	100.00
P5	0.21	0.37	0.00	33.33	50.00
ALL	0.25	0.35	0.01	–	–

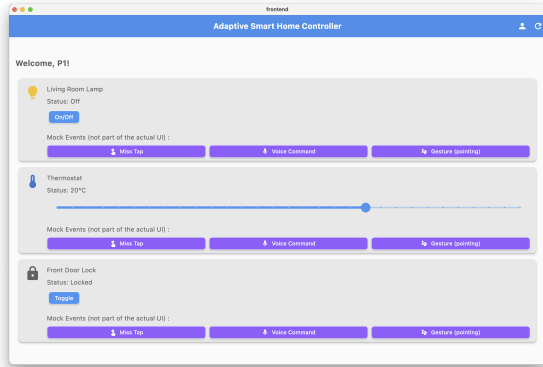
Interpretation. Moving from minimal to maximal rules increases overlap with the LLM (mean Jaccard from 0.25 to 0.35 overall), while exact set matches remain rare (1% overall). This is expected because the LLM often proposes multi-action bundles and tuned parameters that go beyond the canonical rule set. The rule baseline's PAA is highest for the combined needs profile (P4), which aligns with frequent motor and hands free fixes in the event suite; P2 remains at 0% due to the limited visual-focused rules in this baseline. For P3, PAA decreases under the maximal variant because the rule set adds motor fixes that do not align with hands free needs, increasing the denominator without increasing hands free hits.

¹This fast path is a predefined rule based action that can be used instead of the AI logic.

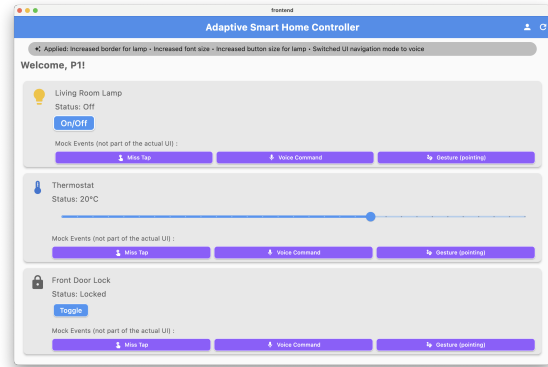
Event counts are identical across profiles (Table 6.1), so differences stem from action selection rather than sample size.

Conclusion. Rules provide useful coverage for obvious corrections and can raise overlap when broadened, but they do not replicate the LLM’s multi-action composition, parameter choices, or personalization. This supports a staged design: route simple patterns through rules for speed, and rely on the LLM and Validator for complex, ambiguous, or profile-sensitive cases. Differences in *rule_PAA_total* across profiles reflect variant choice (minimal vs. maximal) and the top-5 mask rather than unequal event counts.

6.3.13 Micro-cases on P1 and P5 profiles

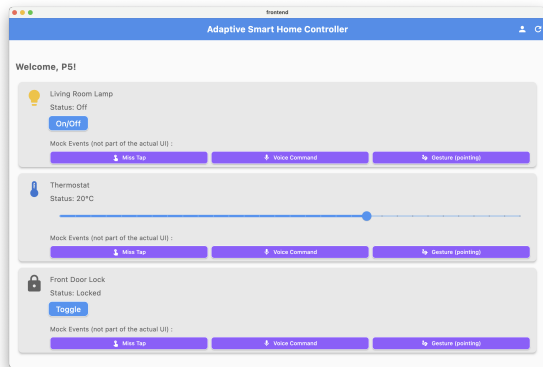


Before (P1, lamp)

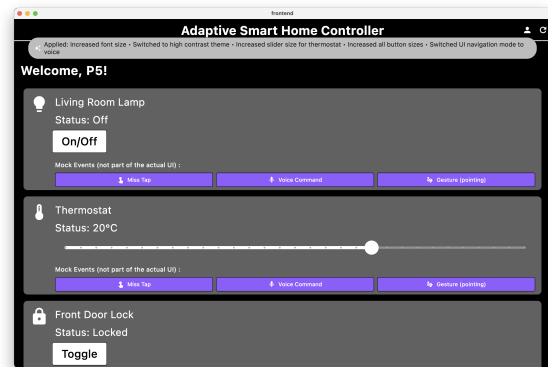


After: larger button and border

Figure 6.1: P1 (motor) miss_tap. The adaptation enlarges the button and border, reducing target difficulty. This aligns with P1’s needs and with the objective scores in Table 6.6 (PAA 50.88%, ERA 57.14%, DCI 0.986).



Before (P5, thermostat)



After: larger slider and higher contrast text

Figure 6.2: P5 (visual+motor) slider_miss. The adaptation increases slider size, text contrast and more as seen by the notification in the app. This matches P5’s mixed needs and the strong scores in Table 6.6 (PAA 77.42%, ERA 64.29%, DCI 1.000).

6.4 Adaptation Performance (Latency)

Latency is reported separately from accessibility and design quality to avoid conflation of speed with effectiveness. The same MA-SIF (balanced) configuration and client-side Δt are used.

6.4.1 Overall summary

Across 84 events and six profiles (Table 6.14):

- **Schema-valid outputs:** 84.52% of events,
- **Classification path:** 100% via Validator,
- **Latency:** p50 13.19 s, p90 17.13 s, max 21.10 s.

Table 6.14: Overall feasibility results with MA-SIF (balanced).

Metric	Value
Events (total)	84
Users (total)	6
Schema-valid outputs (%)	84.52
Validated by Validator (%)	100.00
Latency p_{50} (s)	13.19
Latency p_{90} (s)	17.13
Latency max (s)	21.10

The most frequent actions were:

- `switch_mode` (87): often recommending voice mode.
- `increase_button_size` (74), `increase_button_border` (69), `increase_font_size` (66).
- `increase_slider_size` (28), `increase_contrast` (19).
- Less frequent: `trigger_button` (9), `adjust_spacing` (9).

Because multiple suggestions can be attached to one event, action counts exceed the number of events. Top targets were the lamp, thermostat, and lock, with miss-taps most often on lamp and lock (by design).

Interpretation: The system strongly prioritizes motor-related support (larger targets, borders, spacing) and modal switching to voice when it detects tap/slider errors, exactly the pattern that is needed for motor-impaired and hands-free users. Vision-related support (font size and contrast) also appears consistently, but is less prominent than motor support in this trace.

6.4.2 Per profile outcomes

Validity and latency per profile are in Table 6.15. Latency sits in the 11.8–16.0 s p50 band, with a predictable tail.

Table 6.15: Per profile schema validity and latency under MA-SIF (balanced).

Profile	Declared needs	Valid (%)	p_{50} (s)	p_{90} (s)	max (s)
P0	Baseline	78.57	16.02	20.23	21.10
P1	Motor	85.71	13.19	16.78	17.43
P2	Visual	85.71	13.65	16.46	19.04
P3	Hands free	71.43	11.82	14.68	16.49
P4	Motor + Hands free	85.71	13.63	15.76	17.91
P5	Visual + Motor	100.00	12.70	16.25	16.60

Observations:

- Highest validity appears for combined needs (P5), likely because the event suite provides strong, consistent signals (miss-taps + slider overshoot) that align with the rules and model prior for motor/visual support.
- Lowest validity is P3 (hands-free) at 71.43%. Even though P3 shows the best latency, hands-free preference alone may result in fewer structural UI changes (e.g., fewer size/contrast edits), and the Validator may reject marginal suggestions more often. This indicates room to enrich the hands-free policy (e.g., more explicit voice/gesture affordances and confirmation prompts).
- Latency is in the 11.8–16.0 s p50 band across profiles under the MA-SIF (balanced) config. The spread between p50 and p90 (≈ 3 –4 s) suggests predictable tail behavior, with a single global max near 21 s.

6.4.3 Configuration comparison (all profiles)

For clarity, *Config A* denotes **MA-SIF (balanced)** and *Config B* denotes **MA-SIF (heavy)**. The comparison aggregates all six profiles (P0–P5) across both runs. Table 6.16 reports overall latency and quality, and Table 6.17 breaks results down by profile.

Table 6.16: Overall configuration comparison (lower latency is better; higher Schema-valid, PAA, ERA, and DCI are better).

Config	p50 (s)	Schema-valid (%)	PAA (%)	ERA (%)	DCI
MA-SIF (balanced)	13.19	84.52	55.22	97.62	0.99
MA-SIF (heavy)	36.06	100.00	61.45	98.81	1.00

Table 6.17: Per-profile configuration comparison (p50 latency, Schema-valid, PAA, ERA, and DCI).

Profile	Config	p50 (s)	Schema (%)	PAA (%)	ERA (%)	DCI
P0	MA-SIF (balanced)	16.02	78.57	0.00	100.00	0.98
P0	MA-SIF (heavy)	34.42	100.00	0.00	100.00	0.99
P1	MA-SIF (balanced)	13.19	85.71	50.88	100.00	0.99
P1	MA-SIF (heavy)	34.92	100.00	65.31	100.00	0.99
P2	MA-SIF (balanced)	13.65	85.71	34.69	85.71	1.00
P2	MA-SIF (heavy)	36.55	100.00	40.98	92.86	1.00
P3	MA-SIF (balanced)	11.82	71.43	28.00	100.00	1.00
P3	MA-SIF (heavy)	34.67	100.00	29.17	100.00	1.00
P4	MA-SIF (balanced)	13.63	85.71	80.00	100.00	1.00
P4	MA-SIF (heavy)	40.75	100.00	100.00	100.00	1.00
P5	MA-SIF (balanced)	12.70	100.00	77.42	100.00	1.00
P5	MA-SIF (heavy)	38.54	100.00	77.05	100.00	1.00

Interpretation. MA-SIF (heavy) improves Schema-valid output (100.00% vs 84.52%) and raises PAA (61.45% vs 55.22%), with small gains in ERA and a near-perfect DCI, at the cost of much higher latency (36.06 s vs 13.19 s, about $2.7\times$ slower). Gains are most visible for P1 and P4 where motor and hands free needs dominate; P5 remains similar, and P3 increases only slightly. In practice, the balanced setting is suitable for interactive adaptation, while the heavy setting is better for post-interaction or batch adaptation where stronger alignment and stricter validation justify the extra delay.

6.5 Discussion of Results

Effectiveness. The system focuses almost entirely on accessibility: **97.51%** of all actions are accessibility targeted (Table 6.7). Internal coherence is high across profiles (**DCI \approx 0.995**; Table 6.6). Appropriateness is strong for motor and voice inputs: `miss_tap` and `slider_miss` reach **100%** ERA

and **voice 97.22%**, while **gesture** remains low at **8.33%** (Table 6.9). In the configuration comparison, MA-SIF (heavy) reaches **100.00%** schema-valid outputs and slightly higher ERA and DCI than the balanced setting (Table 6.16), confirming that stricter aggregation and checking can push quality to the top end.

Personalization. Profile–Action Alignment (PAA) is moderate overall ($\sim 55\%$ on top-5 actions) and higher for combined-need profiles (P4, P5), which matches the event suite that emphasises motor errors and mode switching (Table 6.6). MA-SIF (heavy) increases PAA from **55.22%** to **61.45%** overall (Table 6.16) and produces the largest gains where motor and hands free needs dominate (P1, P4; Table 6.17). Some potential improvements include:

1. **Weighted validator scoring:** Rank candidates higher when their category matches the active need, and apply gentle penalties for off-need actions.
2. **Lightweight feedback signals:** Log accept, undo, and reversion in the client and feed these rates back as simple weights to prefer actions users keep.
3. **Richer hands free affordances:** Add confirmation prompts, focus hints, and voice-first navigation so that acceptable sets are available more often, lifting both ERA and PAA for hands free profiles.

Latency. MA-SIF (balanced) achieves a median of **13.19 s** versus **36.06 s** for MA-SIF (heavy) (Table 6.16). The balanced path is more suitable for interactive adaptation, while the heavy path offers stricter outputs for post-interaction or batch contexts. The rule-only sufficiency proxy in Section 6.3.7 and the fast-path analysis indicate that $\sim 86\%$ of events contain at least one corrective action that a simple rule can provide. A staged design is therefore practical: route obvious patterns directly through a rule path and keep the Validator for consolidation, while leaving ambiguous or gesture-driven events to full multi-agent reasoning.

6.6 Study Limitations

- **Feasibility focus.** The evaluation uses structured traces; no human-subject (user) study is included.
- **Trace coverage.** The suite is representative but not exhaustive; real use may show different dwell and timing patterns.
- **Configuration scope.** Two configurations are compared (balanced and heavy). Other variants (for example single-agent SIF or different model budgets) are not explored here.
- **Action menu scope.** The allowed actions form a compact, WCAG-oriented set. This is developer realistic, but it may bias coverage toward needs with generic remedies and does not balance action counts across profiles.
- **Validator transparency.** Rejection reasons are not logged by design, which limits a fine-grained taxonomy of invalid cases; client-side schema checks are used instead.
- **Environment.** Results reflect a single-machine setting with Gemini API calls; other deployments may shift latency.

6.7 Conclusion

The framework delivers valid, accessibility-focused, and coherent adaptations. Objective metrics in Section 6.3 show strong accessibility focus and high internal coherence, with moderate personalization that improves under the heavier configuration. The configuration comparison (Tables 6.16 and 6.17) makes the trade-off explicit: MA-SIF (heavy) raises schema validity and PAA at a substantial latency cost, while MA-SIF (balanced) keeps latency acceptable for asynchronous interaction. The analyses by event type and run stability indicate clear next steps: add need-weighted validator scoring, integrate lightweight user feedback, and extend hands free affordances. These directions increase personalization and keep the schema-based, reproducible evaluation intact, strengthening the case for multimodal, AI-driven adaptation in real-world interfaces.

Chapter 7

Discussion and Future Work

7.1 Overview

This chapter interprets the results of this thesis in the broader context of accessibility and human-computer interaction (HCI), linking the implemented framework to existing research and real deployment. It first discusses implications for adaptive interface design, then summarises key findings and contrasts them with related work. Limitations and threats to validity are addressed explicitly. The chapter closes with a focused future work, spanning near-term engineering improvements and longer-term research directions, including on-device, autonomous adaptation.

7.2 Implications for Accessibility and HCI

Combining deterministic rules with multi-agent, LLM-driven reasoning (SIF/MA-SIF) shows that multi-modal interaction signals can be translated into personalised adaptations during live use. In the feasibility study over **84 events** and **six user profiles**, the system produced **84.5%** schema-valid responses with a **median latency of ~ 13 s**, and all responses were passed through the Validator Agent for consistency. Hands-free profiles trailed slightly in validity, indicating where focused improvements are warranted.

For accessibility, this reduces manual configuration and supports users with motor, visual, or hands-free needs by learning from interaction patterns over time. For HCI, the results indicate a generalisable path towards runtime, cross-platform adaptivity that spans UI, geometry, and input modalities. The hybrid approach also offers a pragmatic design trade-off: fast, predictable fallbacks via rules for common cases; flexible LLM reasoning when context is ambiguous. The caveat is increased complexity in orchestrating two reasoning pipelines and managing latency.

7.3 Key Findings and Contributions

This research contributes to both practical system design and conceptual models for multimodal adaptive UIs:

1. **Framework Architecture:** A modular, three-layer architecture separating input capture, event standardisation, and reasoning, enabling portability across Flutter, SwiftUI, and future platforms (e.g., Unity).
2. **Smart Intent Fusion (SIF):** A hybrid engine combining deterministic rules for guaranteed accessibility fallbacks with LLM-driven reasoning for complex scenarios.
3. **Multi-Agent Extension (MA-SIF):** Specialised agents for UI, geometry, and input, validated by a dedicated Validator Agent to reduce conflicts and hallucinations.
4. **Developer-Focused Integration:** Standardised JSON contracts and generalised methods for event sending, adaptation handling, and profile management.

5. **Accessibility Impact (Empirical):** In the dataset, 97.5% of suggestions directly supported motor, visual, or hands-free needs; median latency remained acceptable for asynchronous adaptation (~ 13 s).

7.4 Comparison with Related Work

SUPPLE [18] SUPPLE demonstrated the value of generating optimised interfaces for specific motor abilities, but relied primarily on constraint-solving over predefined UI descriptions. In contrast, the framework presented in this thesis performs *real-time* adaptation during active use, using multimodal interaction signals fused by a hybrid rule-based and LLM-driven reasoning process. Whereas SUPPLE required a calibration phase, the proposed approach continuously refines adaptations based on interaction history.

Reflow [47] Reflow’s pixel-based UI adaptation enables closed-source applications to be optimised without access to source code, but operates primarily at the visual layout level and does not incorporate higher-level semantic reasoning or multimodal fusion. The framework presented here incorporates profile- and history-driven adaptations across UI, geometry, and input modalities, and supports integration for both new and existing applications through standardised JSON contracts.

UICoder [47] UICoder automates the generation of UI code from textual descriptions, enabling rapid interface creation. The approach taken in this work is complementary rather than competitive: whereas UICoder generates new interfaces at compile time, the proposed framework adapts *existing* live interfaces at runtime, guided by multimodal context and inferred user intent.

GUIDe [23] GUIDe’s gaze-augmented interaction techniques illustrate the potential of combining modalities for improved accuracy and reduced activation errors. This framework generalises that principle by supporting arbitrary modality combinations (e.g., voice + touch, gaze + gesture) and applying them within a cross-platform adaptation pipeline.

Reinforcement Learning-based UI Adaptation [20] Reinforcement learning approaches can personalise UIs over time using physiological feedback, but often require specialised sensing hardware and extended training periods. The framework presented in this thesis prioritises lightweight integration and immediate adaptation using LLM reasoning, while retaining the potential to incorporate reinforcement learning techniques in future iterations.

Table 7.1 summarises key differences between the framework and representative related work.

Table 7.1: Comparison of related adaptive UI systems

System	Modalities supported	Supported	Reasoning Method	Adaptation Scope	Extensibility
SUPPLE	Single modality (motor calibration)		Constraint optimisation	Layout generation (compile-time)	Requires redesign for new modalities
Reflow	Any (pixel-based only)		Visual analysis + rules	Layout changes only (runtime)	Limited; no semantic adaptation
UICoder	N/A (code generation)		LLM code synthesis	New UI creation (compile-time)	Extendable via prompt tuning
GUIDe	Gaze + keyboard/-mouse		Rule-based	Targeting and scrolling (runtime)	Limited to gaze/-pointing
RL-based UI Adaptation	Depends on sensors		Reinforcement learning	Layout, content (runtime)	Requires specialised hardware
This work	Touch, voice, gesture, keyboard (+ extensible)		Hybrid rules + multi-agent LLM	UI, geometry, input modes (runtime)	High; JSON contracts + config-based agents

7.5 Future Work

While the framework in its current first iteration form demonstrates the feasibility of (close to) real-time, multimodal adaptive UI adaptation, several opportunities exist to expand its capabilities, improve its performance, and extend its reach into new domains. Future work can be divided into short-term engineering improvements (such as SDK packaging and support for additional modalities) and long-term research directions (such as autonomous, on-device reasoning agents capable of compile-time and runtime adaptation).

7.5.1 Short-Term Improvements

Improved Gesture Support

The evaluation revealed that gesture events achieved only **8.33%** ERA (Table 6.9), significantly lower than touch-based interactions (100%) and voice commands (97.22%). This limitation stems from insufficient disambiguation mechanisms and limited gesture-specific adaptations in the current action set.

To address this, future iterations should implement a gesture disambiguation pipeline that includes brief confirmation prompts (e.g., "Point to confirm?") and fallback routing to voice modality when gesture recognition confidence is low. Additionally, extending the acceptable action sets with gesture-specific affordances such as focus outlines, dwell-to-activate mechanisms, and spatial targeting hints would improve appropriateness.

Enhanced Personalization for Single-Need Profiles

While combined-need profiles (P4, P5) achieved strong Profile-Action Alignment scores, single-need profiles showed moderate performance, with P2 and P3 achieving 34.69% and 28.00% PAA respectively (Table 6.6). This indicates room for improvement in tailoring adaptations to specific accessibility needs.

The solution involves implementing need-weighted scoring within the Validator Agent, where adaptations matching the active profile's declared needs receive higher priority. The rule fallback engine could be implemented in a similar way for a higher overall effectiveness. Additionally, integrating lightweight user feedback mechanisms (accept, undo, revert actions) would provide continuous learning signals to refine personalization over time. These improvements should target a **+10 point** increase in PAA for underperforming profiles while maintaining high Design Coherence Index ($DCI \geq 0.99$).

Latency Reduction for Interactive Use

Current median latency of **13.19s** for MA-SIF (balanced) configuration, while acceptable for asynchronous adaptation, could benefit from optimisation for more responsive interaction. The heavy configuration's 36.06s median latency further emphasizes the need for speed improvements.

A staged approach would route obvious patterns (miss-tap \rightarrow larger target, slider overshoot \rightarrow larger slider) through a fast rule-based path, followed by MA-SIF consolidation for complex cases. Caching frequently suggested adaptations and implementing selective on-device classification for path selection could further reduce response times. The target is achieving $p50 \leq 8s$ while maintaining current schema validity and ERA performance.

SDK Development and Cross-Platform Support

To facilitate adoption, the framework requires packaging as developer-friendly SDKs for Flutter, SwiftUI, Unity/VR, and web platforms. These SDKs should include pre-built adapter components, standardised integration patterns, debugging tools, and comprehensive documentation with accessibility-focused examples.

The SDK should also support offline modes using lightweight, on-device models for basic adaptations when cloud connectivity is unavailable. This addresses both privacy concerns and deployment constraints in sensitive environments.

7.5.2 Research Directions

Visual UI Understanding and Semantic Analysis

Current reasoning relies on developer-provided metadata about UI elements and their relationships. Integrating computer vision capabilities would enable the system to analyse live interface screenshots, understanding spatial layouts, colour schemes, element hierarchies, and accessibility violations without requiring manual annotation.

This visual understanding could be implemented using fine-tuned vision transformers or CLIP-like models trained on UI datasets. The analyser would complement existing metadata with real-time visual context, enabling more nuanced adaptations such as repositioning elements based on visual density or adjusting contrast based on actual colour relationships.

Specialised AI Models for UI Adaptation

Moving beyond general-purpose LLMs towards domain-specific models trained on UI interaction logs, accessibility guidelines (WCAG), and adaptation effectiveness data would improve both accuracy and efficiency. These models could be fine-tuned for specific contexts (healthcare, gaming, productivity) and optimised for on-device deployment.

Such specialisation would reduce dependence on external APIs, improve privacy, and enable faster response times. The models could incorporate reinforcement learning from user feedback to continuously improve adaptation strategies.

Comprehensive User Studies

While the feasibility study demonstrates technical capability, validating real-world effectiveness requires comprehensive user studies with participants who have actual accessibility needs. A properly powered study should recruit N=12-18 participants across three groups: motor-impaired (4-6 participants), visually impaired (4-6 participants), and hands-free users (4-6 participants).

The study design should employ a within-subjects comparison, testing participants on both rule-based adaptations and the full MA-SIF pipeline using realistic smart home control scenarios. Objective measures would include task completion time, error rates, and adaptation acceptance rates, while subjective measures would capture perceived usefulness through validated instruments such as the System Usability Scale (SUS) and NASA Task Load Index (TLX).

Critically, participants should include individuals with lived experience of the target accessibility constraints rather than simulated impairments. This ensures that findings reflect genuine usability improvements and adaptation preferences. Such studies would provide essential evidence for the framework's practical impact beyond technical feasibility, while identifying refinements needed for real-world deployment based on authentic user experiences.

Extended Modality Support

Future work should expand input modality support to include eye tracking for hands-free navigation, advanced gesture recognition using depth sensors or computer vision, and brain-computer interfaces for users with severe motor impairments. Each new modality requires careful integration with the existing event standardisation pipeline and adaptation logic.

Particular attention should be paid to multimodal fusion strategies that can intelligently combine signals from multiple simultaneous inputs, handling confidence weighting and temporal alignment to produce coherent user intent inference.

Privacy-Preserving and Federated Learning

As the framework collects rich interaction data including multimodal inputs and adaptation effectiveness, implementing privacy-preserving techniques becomes crucial. Federated learning approaches could enable model improvement across users without centralising sensitive data, while differential privacy techniques could protect individual interaction patterns. On-device processing capabilities should be expanded to minimise data transmission requirements, with selective cloud processing only for complex reasoning tasks that exceed local computational capacity.

Chapter 8

Conclusion

8.1 Summary of the Work

This thesis designed and evaluated a multimodal AI-driven GUI framework that adapts user interfaces in near real-time. The goal was to support personalized and accessible experiences, especially in health contexts where users may have motor, visual, or hands free constraints. The framework captures touch, keyboard, voice, and gesture inputs and turns them into concrete adaptation actions via *Smart Intent Fusion* (SIF).

The architecture has three modular layers: an **Input Adapter** that standardises events and handles basic management, a **SIF Backend** that combines rules with LLM-driven reasoning, and a **Frontend** that renders the UI and applies adaptations. A key addition is **Multi-Agent SIF** (MA-SIF): separate specialists for UI, geometry, and input, with a Validator Agent to reconcile and validate outputs. This setup reduced hallucinations, improved reliability, and stayed easy to configure through an external JSON file. SIF first started from a single-agent LLM and moved to MA-SIF after seeing prompt creep and uneven JSON adherence; splitting roles plus validation made outputs more predictable without sacrificing adaptability.

Research Objectives and Outcomes To reach the overall goal, six concrete objectives were set. Each is briefly restated with its outcome.

1. **Develop a modular, cross-platform architecture.** A three-layer framework (Frontend, Input Adapter, Backend) was implemented with a shared JSON-based event and adaptation format. It integrates with Flutter today and keeps a clear path for SwiftUI and Unity through the same contracts.
2. **Implement multimodal input fusion.** Touch, keyboard, voice, and gesture inputs are supported and logged in a unified event schema. The pipeline fuses these with profile data and history, enabling context-aware adaptations across modalities.
3. **Integrate Smart Intent Fusion (SIF).** A hybrid engine was built that blends deterministic rules with multi-agent LLM reasoning (MA-SIF). The Validator Agent enforces schema, removes duplicates, and prevents conflicts, which yields high internal coherence (DCI \approx **0.995**; Table 6.6).
4. **Deliver accessibility-focused adaptations.** The system consistently proposes accessibility actions: **97.51%** of all suggestions target motor, visual, or hands free needs (Table 6.7). Typical changes include enlarging targets, increasing contrast and font size, and switching interaction modes.
5. **Provide a developer-friendly integration path.** The framework exposes simple contracts and reusable patterns; adapters apply changes on the client with minimal glue code. Configuration lives in JSON and can be swapped without redeploying the app.
6. **Evaluate system performance and impact.** A feasibility study over **84 events** across **six profiles** assessed adaptation quality and latency. In the balanced configuration, median latency is **13.19 s**, schema-valid outputs **84.52%**, and profile-action alignment (top five actions) **55.22%**.

A heavier configuration raises schema-valid to **100.00%** and PAA to **61.45%**, at **36.06 s** median latency (Tables 6.16–6.17).

8.2 Summary of Contributions

1. **Modular, cross-platform architecture:** Designed to work across Flutter and SwiftUI, with a clear path to future platforms such as Unity/VR through shared JSON contracts.
2. **Smart Intent Fusion (SIF):** A hybrid engine that mixes deterministic rules with multi-agent LLM reasoning for ambiguous cases.
3. **MA-SIF extension:** Specialist agents (UI/geometry/input) plus a Validator Agent for conflict resolution and schema adherence, yielding high coherence ($DCI \approx 0.995$).
4. **Developer-friendly integration:** Consistent contracts for events, adaptations, and profiles; minimal glue code; configuration in JSON.
5. **Accessibility focus (empirical):** **97.51%** of actions are accessibility targeted; strong performance on common motor/voice issues and moderate personalization that improves under the heavy configuration.

8.3 Limitations

Current limitations include:

- **LLM API dependency:** Cloud models require connectivity and are subject to quota and availability.
- **Synthetic evaluation:** Results come from structured profiles and tasks; a user study is still needed to capture acceptance and perceived control.
- **UI context awareness:** Reasoning relies on predefined metadata; there is no live semantic understanding of layouts yet.
- **Latency:** Multi-agent LLM reasoning adds seconds of delay, which may not suit highly time-critical flows.
- **Privacy and cost:** Multimodal logging can be sensitive and cloud inference can be costly; production systems need strict telemetry controls and efficiency measures.

These constraints make the current prototype best suited for exploratory deployments rather than safety-critical settings. Dependence on a cloud LLM introduces availability risks; if the model stalls or rate-limits, behaviour regresses to rules. The evaluation, while systematic, is based on scripted scenarios and may miss edge cases. Without a live UI model, layout nuances can be overlooked. Latency is noticeable for highly interactive flows. Finally, rich logging raises privacy and cost concerns. Near-term mitigations include progressive enhancement (rules first, MA-SIF second), batching and caching, rate-aware fallbacks, on-device redaction, and a UI analyser to bridge the context gap.

8.4 Lessons Learned

- **Hybrid > pure AI.** Rules keep the interface responsive; MA-SIF adds smarter, profile-aware changes moments later.
- **Schemas are leverage.** Stable JSON contracts made the system portable and easy to debug as components evolved.
- **Prompts and validation matter.** Constraint-focused prompts plus a Validator Agent improved output quality and reduced contradictions.
- **Simulation helps.** Synthetic profiles and events were enough to tune latency limits, prompts, and fallback behaviour before involving real users.
- **Modularity reduces risk.** Loose coupling made it straightforward to swap models and modalities and points to a clear path toward UI-aware vision and on-device reasoning.

8.5 Evaluation Results: Conclusion

The Adaptive Smart Home Controller shows that the framework can deliver valid, accessibility-focused adaptations in near real-time. Across **84 events** and six profiles, the balanced configuration reached **84.52%** schema-valid responses (i.e., final payloads that pass the strict JSON schema), with **100%** of responses flowing through the Validator Agent. **97.51%** of suggested actions were accessibility targeted (Table 6.7), and internal coherence was high (DCI \approx **0.995**; Table 6.6).

Effectiveness is strong on common motor and voice cases: *missed tap* and *slider overshoot* reached **100%** ERA, and *voice command* **97.22%**; gestures remain the main opportunity for improvement (**8.33%**; Table 6.9). Personalization is moderate overall (PAA **55.22%** on the top five actions) and highest for combined needs (P4, P5), which aligns with the action mix. Latency under MA-SIF (balanced) is **13.19 s** median (p90 **17.13 s**, max **21.10 s**), which suits asynchronous adaptation. The heavier configuration increases schema validity to **100.00%** and raises alignment to **61.45%** at **36.06 s** median latency (Tables 6.16–6.17), making the trade-off explicit.

A rules-only engine did *not* outperform this setup. Overlap with the final LLM+Validator sets was limited (Jaccard \sim 0.25–0.35 overall; exact matches \sim 1%), and rules missed multi-action bundles and need-aware choices (Table 6.13). That said, a rule *fast path* is still useful: in about **86%** of events there is at least one canonical corrective action available (aggregated from Table 6.9), so emitting that immediately and letting the Validator consolidate afterwards can lower median latency without sacrificing quality.

Overall, the combination of rules + MA-SIF + validation produced predictable outputs without sacrificing adaptability. The next step is a small, focused user study to measure acceptance and perceived control, and to validate that the balanced configuration offers the right quality–latency point in practice. Beyond that, UI-aware action variants and light feedback signals are expected to lift personalization for single-need profiles while keeping coherence and latency within bounds.

8.6 Self-Reflection and Reflection on the Work

This project taught me how to turn an ambitious idea about multimodal adaptive interfaces into a working system through many small steps. I learned that a hybrid approach works best. Simple rules keep the interface responsive, while LLM reasoning adds context aware changes. Stable JSON contracts were a key part of the design because they gave every layer a shared language and made testing and extension much easier. I also learned to write clear prompts and to validate every result with a strict schema so that the output stays reliable and useful.

I solved core challenges such as standardising events across input types, using WebSocket for fast feedback and HTTP for profile management, and reducing hallucinations with a multi agent setup and a strong Validator Agent. Some things were out of scope for now, such as live semantic understanding of the UI, very low latency on device models, and large user studies. Because of these limits the system is ready for exploration and demos rather than strict time critical or safety critical use.

This work also provided me with insights into future directions. A UI analyser that uses visual input or a widget tree would let the backend reason over the real screen. Light domain specific models that run on device could lower latency a lot. Studies with people who have motor or visual impairments would show real effects and acceptance. Privacy by design logging and redaction at the edge would make the system safer for sensitive data.

My process improved as I narrowed the scope. By iterating on a small end to end pipeline from adapter to backend to frontend I got faster feedback and clearer priorities. Simulated profiles and scripted events were enough to tune latency budgets, prompts, and fallbacks before involving real users. I also made progress measurable with schema checks, logging, and small comparisons instead of guesswork.

I learned that I work well between research and engineering. I enjoyed turning ideas like MA-SIF and its agents into robust code. I also learned to avoid trying to do everything at once. A focused first version and clear contracts helped me move fast without losing quality. I became more precise in writing prompts, planning for failure with timeouts and fallbacks, and stating privacy and cost assumptions. I also learned to accept uncertainty since LLM systems do not always behave in predictable ways. By evaluating the system more and more, I came to the conclusion that this has real potential for improving user experience.

Looking back, the main aims were met. The system delivers near real-time adaptations, fuses multimodal intent, and stays friendly for developers. Rules give quick feedback and MA-SIF adds richer changes. The main gaps are latency and deep context awareness. These are not failures but a roadmap. On device inference, semantic UI analysis, and real world studies are the next steps.

8.7 Final Remarks

This thesis shows that **Smart Intent Fusion**, in a multi-agent setup with a Validator, delivers *reliable, accessibility-focused* adaptation at interactive pace. In the six-profile, 84-event study, almost all suggestions targeted accessibility goals (97.51%), internal coherence was high ($DCI \approx 0.995$), and the balanced configuration kept median latency near 13s. Effectiveness was near-perfect on common motor/voice cases (missed tap, slider overshoot, voice command), with gestures marked as the main opportunity for refinement.

A key outcome is that a *rules-only* policy did not outperform this architecture. Overlap with the final LLM+Validator sets stayed modest (Jaccard ~ 0.25 – 0.35 ; exact matches $\sim 1\%$), and rules missed multi-action bundles and need-aware choices that lift personalization. At the same time, the data support a practical compromise: in $\sim 86\%$ of events there is at least one canonical rule fix available, so a *fast path* can emit that immediately while MA-SIF and the Validator consolidate the full, profile-aware adaptation a moment later. In short, rules help with speed; the hybrid pipeline wins on coverage, composition, and alignment.

Looking ahead, there is some work to do: package the adapter and contracts as SDKs (Flutter, SwiftUI, future Unity/VR), add offline and privacy-preserving modes, and ground decisions in real UI context via widget trees and vision. A small user study with motor-, visual-, and hands free participants will turn feasibility into user-level evidence (acceptance, perceived control). Beyond that, UI-aware action variants and light feedback signals can raise personalization for single-need profiles without sacrificing coherence or latency.

Treat UIs as programmable surfaces, fuse multimodal intent, and validate before acting. With that recipe, interfaces move from static screens to adaptive systems that are personalized, inclusive, and resilient to the diversity of human interaction.

References

- [1] Pierre A. Akiki, Arosha K. Bandara, and Yijun Yu. “Integrating adaptive user interface capabilities in enterprise applications”. In: *Companion Proceedings of the 36th International Conference on Software Engineering (ICSE)*. ACM, 2014, pp. 487–490. DOI: 10.1145/2568225.2568230.
- [2] Pradeep Atrey et al. “Multimodal fusion for multimedia analysis: A survey”. In: *Multimedia Syst.* 16 (Nov. 2010), pp. 345–379. DOI: 10.1007/s00530-010-0182-0.
- [3] Greg J. Badros, Alan Borning, and Peter J. Stuckey. “The Cassowary Linear Arithmetic Constraint Solving Algorithm”. In: *ACM Transactions on Computer-Human Interaction* 8.4 (2001), pp. 267–306. DOI: 10.1145/504704.504705.
- [4] Matthias Baldauf, Schahram Dustdar, and Florian Rosenberg. “A survey on context-aware systems”. In: *International Journal of Ad Hoc and Ubiquitous Computing* 2.4 (2007), pp. 263–277. DOI: 10.1504/IJAHUC.2007.014070.
- [5] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. *Multimodal Machine Learning: A Survey and Taxonomy*. 2017. arXiv: 1705.09406 [cs.LG]. URL: <https://arxiv.org/abs/1705.09406>.
- [6] Richard A. Bolt. ““Put-that-there”: Voice and gesture at the graphics interface”. In: *Proceedings of the 7th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH ’80. Seattle, Washington, USA: Association for Computing Machinery, 1980, pp. 262–270. ISBN: 0897910214. DOI: 10.1145/800250.807503. URL: <https://doi.org/10.1145/800250.807503>.
- [7] Andreas Bulling et al. “Gaze interaction in the post-WIMP world”. In: May 2012, pp. 1221–1224. ISBN: 9781450310161. DOI: 10.1145/2212360.2212428.
- [8] Gaëlle Calvary et al. “A Unifying Reference Framework for Multi-Target User Interfaces”. In: *Interacting with Computers* 15.3 (2003), pp. 289–308. DOI: 10.1016/S0953-5438(03)00010-9.
- [9] Ananya Choudhury, Anjan Talukdar, and Kandarpa Sarma. “A Review on Vision-Based Hand Gesture Recognition and Applications”. In: Aug. 2015, pp. 261–286. DOI: 10.4018/978-1-4666-8493-5.ch011.
- [10] Andy Cockburn, Carl Gutwin, and Jason Alexander. “Improving Pointing Performance by Mixing Control and Display Space: Control-Display Gain and Magnetic Pointer Fences”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Florence, Italy: ACM, 2008, pp. 129–138. DOI: 10.1145/1357054.1357074.
- [11] Anind K. Dey. “Understanding and using context”. In: *Personal and Ubiquitous Computing* 5.1 (2001), pp. 4–7. DOI: 10.1007/s007790170019. URL: <https://doi.org/10.1007/s007790170019>.
- [12] Morgan Dixon, James Fogarty, and Jacob Wobbrock. “A general-purpose target-aware pointing enhancement using pixel-level analysis of graphical interfaces”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’12. Austin, Texas, USA: Association for Computing Machinery, 2012, pp. 3167–3176. ISBN: 9781450310154. DOI: 10.1145/2207676.2208734. URL: <https://doi.org/10.1145/2207676.2208734>.
- [13] Andrew T. Duchowski. *Eye Tracking Methodology: Theory and Practice*. 3rd ed. eBook ISBN: 978-3-319-57883-5, Published: 24 May 2017, Softcover ISBN: 978-3-319-57881-1, Published: 07 June 2017, Number of Pages: XL, 366, Number of Illustrations: 71 b/w illustrations, 62 illustrations in colour. Cham: Springer Cham, 2017. ISBN: 978-3-319-57881-1. DOI: 10.1007/978-3-319-57883-5. URL: <https://link.springer.com/book/10.1007/978-3-319-57883-5>.
- [14] Leah Findlater and Joanna McGrenere. “A Comparison of Static, Adaptive, and Adaptable Menus”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’04)*. ACM, 2004, pp. 89–96. DOI: 10.1145/985692.985704.

- [15] Leah Findlater et al. “Ephemeral Adaptation: The Use of Gradual Onset to Improve Menu Selection Performance”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, 2009, pp. 1655–1664. DOI: 10.1145/1518701.1518956.
- [16] Paul M. Fitts. “The information capacity of the human motor system in controlling the amplitude of movement”. In: *Journal of Experimental Psychology* 47.6 (1954), p. 381. DOI: 10.1037//0096-3445.121.3.262.
- [17] Kathleen Kara Fitzpatrick, Alison Darcy, and Molly Vierhile. “Delivering Cognitive Behavior Therapy to Young Adults With Symptoms of Depression and Anxiety Using a Fully Automated Conversational Agent (Woebot): A Randomized Controlled Trial”. In: *JMIR Ment Health* 4.2 (June 2017), e19. ISSN: 2368-7959. DOI: 10.2196/mental.7785. URL: <http://www.ncbi.nlm.nih.gov/pubmed/28588005>.
- [18] Krzysztof Gajos and Daniel S. Weld. “SUPPLE: automatically generating user interfaces”. In: *Proceedings of the 9th International Conference on Intelligent User Interfaces*. IUI '04. Funchal, Madeira, Portugal: Association for Computing Machinery, 2004, pp. 93–100. ISBN: 1581138156. DOI: 10.1145/964442.964461. URL: <https://doi.org/10.1145/964442.964461>.
- [19] Krzysztof Z. Gajos, Daniel S. Weld, and Jacob O. Wobbrock. “Automatically Generating Personalized User Interfaces with SUPPLE”. In: *Artificial Intelligence* 174.12–13 (2010), pp. 910–950. DOI: 10.1016/j.artint.2010.05.005.
- [20] Daniel Gaspar-Figueiredo. *Learning from Interaction: User Interface Adaptation using Reinforcement Learning*. 2023. arXiv: 2312.07216 [cs.SE]. URL: <https://arxiv.org/abs/2312.07216>.
- [21] Matthew B. Hoy. “Alexa, Siri, Cortana, and More: An Introduction to Voice Assistants”. In: *Medical Reference Services Quarterly* 37.1 (2018). PMID: 29327988, pp. 81–88. DOI: 10.1080/02763869.2018.1404391. eprint: <https://doi.org/10.1080/02763869.2018.1404391>. URL: <https://doi.org/10.1080/02763869.2018.1404391>.
- [22] Anthony Jameson. “Adaptive Interfaces and Agents”. In: *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*. Ed. by Julie A. Jacko and Andrew Sears. Lawrence Erlbaum Associates, 2003, pp. 305–330. ISBN: 0805838384.
- [23] Manu Kumar and Terry Winograd. “GUIDe: gaze-enhanced UI design”. In: *CHI '07 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '07. San Jose, CA, USA: Association for Computing Machinery, 2007, pp. 1977–1982. ISBN: 9781595936424. DOI: 10.1145/1240866.1240935. URL: <https://doi.org/10.1145/1240866.1240935>.
- [24] Manu Kumar and Terry Winograd. “GUIDe: gaze-enhanced UI design”. In: *GUIDe: Gaze-enhanced UI Design, CHI: Conference on Human Factors in Computing Systems*. Apr. 2007, pp. 1977–1982. DOI: 10.1145/1240866.1240935.
- [25] Håkon Wium Lie et al. *Media Queries Level 3*. Tech. rep. W3C Recommendation. World Wide Web Consortium (W3C), 2012. URL: <https://www.w3.org/TR/mediaqueries-3/>.
- [26] Yi-Jung Lin and Stuart Speedie. “Role-Based and Adaptive User Interface Designs in a Teledermatology Consult System: A Way to Secure and a Way to Enhance”. In: *AMIA Annual Symposium Proceedings*. American Medical Informatics Association, 2003, p. 913. URL: <https://pmc.ncbi.nlm.nih.gov/articles/PMC1480225/>.
- [27] Camillo Lugaresi et al. *MediaPipe: A Framework for Building Perception Pipelines*. 2019. arXiv: 1906.08172 [cs.DC]. URL: <https://arxiv.org/abs/1906.08172>.
- [28] Kris Luyten and Karin Coninx. “An XML-Based Runtime User Interface Description Language for Mobile Computing Devices”. In: *Interactive Systems: Design, Specification, and Verification*. Ed. by Chris Johnson. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001, pp. 1–15. DOI: 10.1007/3-540-44664-0_1. URL: https://doi.org/10.1007/3-540-44664-0_1.
- [29] Kris Luyten, Kristof Thys, and Karin Coninx. “Profile-Aware Multi-Device Interfaces: An MPEG-21-Based Approach for Accessible User Interfaces”. In: (2005). DOI: 10.14236/ewic/AD2005.5.
- [30] Ethan Marcotte. *Responsive Web Design*. A List Apart, Issue 306. 2010. URL: <https://alistapart.com/article/responsive-web-design/>.
- [31] Ethan Marcotte. “Responsive Web Design”. In: *A List Apart* 306 (2010). URL: <https://alistapart.com/article/responsive-web-design>.
- [32] Jerzy Mizeraczyk et al. “A Method for Underwater Wireless Data Transmission in a Hydroacoustic Channel under NLOS Conditions”. In: *Sensors* 21.23 (2021). ISSN: 1424-8220. DOI: 10.3390/s21237825. URL: <https://www.mdpi.com/1424-8220/21/23/7825>.
- [33] Florian Mueller and Damon Young. “10 Lenses to Design Sports-HCI”. In: *Foundations and Trends® in Human-Computer Interaction* 12.3 (2018), pp. 172–237. ISSN: 1551-3955. DOI: 10.1561/11000000076. URL: <http://dx.doi.org/10.1561/11000000076>.

- [34] Brad A. Myers, Scott E. Hudson, and Randy Pausch. “Past, present, and future of user interface software tools”. In: *ACM Transactions on Computer-Human Interaction* 7.1 (2000), pp. 3–28. DOI: 10.1145/344949.344959.
- [35] Francisco R. Ortega et al. “Gesture elicitation for 3D travel via multi-touch and mid-Air systems for procedurally generated pseudo-universe”. In: *2017 IEEE Symposium on 3D User Interfaces (3DUI)*. 2017, pp. 144–153. DOI: 10.1109/3DUI.2017.7893331.
- [36] Sharon Oviatt. “Ten myths of multimodal interaction”. In: *Commun. ACM* 42.11 (Nov. 1999), pp. 74–81. ISSN: 0001-0782. DOI: 10.1145/319382.319398. URL: <https://doi.org/10.1145/319382.319398>.
- [37] Sharon Oviatt, Rachel Coulston, and Rebecca Lunsford. “When do we interact multimodally? cognitive load and multimodal communication patterns”. In: *Proceedings of the 6th International Conference on Multimodal Interfaces*. ICMI ’04. State College, PA, USA: Association for Computing Machinery, 2004, pp. 129–136. ISBN: 1581139950. DOI: 10.1145/1027933.1027957. URL: <https://doi.org/10.1145/1027933.1027957>.
- [38] Zhanna Sarsenbayeva et al. “Methodological Standards in Accessibility Research on Motor Impairments: A Survey”. In: *ACM Computing Surveys* 55.7 (2022), pp. 1–35. DOI: 10.1145/3543509.
- [39] Bill N. Schilit and Marvin M. Theimer. “Context-aware computing applications”. In: *IEEE Workshop on Mobile Computing Systems and Applications* (1994), pp. 85–90. DOI: 10.1109/WMCSA.1994.16.
- [40] Andrew Sears and Ben Shneiderman. “Split Menus: Effectively Using Selection Frequency to Organize Menus”. In: *ACM Transactions on Computer-Human Interaction* 1.1 (1994), pp. 27–51. DOI: 10.1145/174630.174632.
- [41] Dídac Surís, Sachit Menon, and Carl Vondrick. *ViperGPT: Visual Inference via Python Execution for Reasoning*. 2023. arXiv: 2303.08128 [cs.CV]. URL: <https://arxiv.org/abs/2303.08128>.
- [42] T. R. Toelle et al. “App-based multidisciplinary back pain treatment versus combined physiotherapy plus online education: a randomized controlled trial”. In: *NPJ Digital Medicine* 2.1 (May 2019), p. 34. DOI: 10.1038/s41746-019-0109-x.
- [43] Bryan Wang, Gang Li, and Yang Li. “Enabling Conversational Interaction with Mobile UI using Large Language Models”. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. CHI ’23. Hamburg, Germany: Association for Computing Machinery, 2023. ISBN: 9781450394215. DOI: 10.1145/3544548.3580895. URL: <https://doi.org/10.1145/3544548.3580895>.
- [44] Daniel Wigdor and Dennis Wixon. *Brave NUI World: Designing Natural User Interfaces for Touch and Gesture*. 1st. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2011. ISBN: 0123822319.
- [45] Jacob O. Wobbrock et al. “Ability-Based Design: Concept, Principles, and Examples”. In: *ACM Transactions on Accessible Computing* 3.3 (2011). DOI: 10.1145/1952383.1952384.
- [46] Jacob O. Wobbrock et al. “The Angle Mouse: Target-Agnostic Dynamic Gain Adjustment Based on Angular Deviation”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’09)*. ACM, 2009, pp. 1401–1410. DOI: 10.1145/1518701.1518912.
- [47] Jason Wu. “Computational Understanding of User Interfaces”. In: (Aug. 2024). DOI: 10.1184/R1/26485132.v1. URL: https://kilthub.cmu.edu/articles/thesis/Computational_Understanding_of_User_Interfaces/26485132.
- [48] Shunyu Yao et al. *ReAct: Synergizing Reasoning and Acting in Language Models*. 2023. arXiv: 2210.03629 [cs.CL]. URL: <https://arxiv.org/abs/2210.03629>.
- [49] Boyuan Zheng et al. *GPT-4V(ision) is a Generalist Web Agent, if Grounded*. 2024. arXiv: 2401.01614 [cs.IR]. URL: <https://arxiv.org/abs/2401.01614>.