GENERATIVE AI & UX

FROM LINEAR JOURNEYS TO TAILORED TERRAIN

MARC SEEFELDER, MILES R JOHNSON

SUMMARY

In today's swiftly evolving digital realm, we examine the monumental shift towards hyper-personalized and dynamic user experiences, all driven by the unprecedented capabilities of generative artificial intelligence. Our exploration focuses on the strategic convergence of GenAI and advanced UX design, illuminating a roadmap crucial for businesses aiming to stay ahead. This paper presents a comprehensive analysis, revealing the imminent challenges and unparalleled opportunities that await.

MING LABS

GENAI & UX
CONTENTS

INTRODUCTION

In the history of computing, we are at a pivotal moment. We've moved from the early days of batch processing to 60 years of command-based interactions. Now, as highlighted by Jakob Nielsen in his piece, AI: First New UI Paradigm in 60 Years, we're transitioning to a new phase: intent-based user interactions, marking a significant evolution in how we interact with technology.

But what does this shift in paradigm imply for user experience (UX) design? How do we reimagine and craft digital interactions for this new paradigm? This paper seeks to answer that very question, charting the progression from linear UX designs to the emergent landscape of dynamic hyper-personalised experiences. What are the challenges, opportunities, and ultimate implications of this seismic shift in design thinking? Join us in exploring the future of UX in the age of intent-based interactions.

At the intersection of advanced artificial intelligence (GenAI) and state-of-the-art UX design, this paper examines the transformative horizon of digital experiences. Moving beyond the fixed contours of linear paths, we envision a digital realm that intuitively reshapes itself to users' desires. While we draw inspiration from thought leaders like Nielsen, our exploration is equally grounded in hands-on experiences from two recent project contexts at the convergence of GenAI and UX, music licensing for content creators, and old-age care support. Central to our discourse is the pivotal question: How can we transition from static journeys to dynamic hyper-personalised interactions? As we delve into pivotal cases and examine emerging design methodologies, our vision is a future

(HYPER)PERSONALISATION

In today's rapidly evolving digital ecosystem, a compelling narrative underscores the importance of UX: the shift towards bespoke personalisation. Modern users no longer resonate with generic interfaces; they seek experiences intricately woven around their unique preferences and behaviours. It's hard to pinpoint where this desire came from, but it's not haphazard to guess that some of this desire for personalisation has been learned from our content algorithms, prevalent across news, social media, and e-commerce. So, if content recommendations can be managed via an algorithm – why not our user journeys, why not our interfaces, why not everything surrounding our digital products?

UNPACKING PERSONALISATION

At its core, personalisation in UX signifies adapting various platform elements to cater to individual user nuances. This adaptation manifests in several ways, from content to tone, to user journeys and layouts. Personalised content refers to the actual text or media presented to the user. This could be as simple as using a user's name in greetings or as complex as adjusting articles based on age, reading ability, language, or cultural background. Tone refers to the mood or sentiment of the content. A platform might adopt a formal tone for professional users while switching to a casual, friendly tone for younger users. User Journeys relate to the sequence or pathway a user follows within a platform. For instance, a first-time user might be shown only core services initially, with advanced features surfacing over time based on the user's objectives. Finally, layouts are the visual arrangement of elements on a page. Depending on user preferences or their device, the layout might be grid-based, list-based, dark mode, light mode, and more.

WHY PERSONALISATION MATTERS IN DESIGN

Personalising these aspects without the product experience has at the best of times been approached with a research lens, using a combination of background research and contextual user interviews to generate personas and their corresponding desired experience. These personas are often prioritised and user stories are then applied to each (with core stories shared across all types, such as account validation). From a marketing perspective, we would call this segmentation – separating our users based on demographics, needs, and desires. This approach continues to be an extremely effective tool for product teams in breaking down larger service offerings into pieces that could be defined, designed, built, tested, and deployed with confidence.

WHY PERSONALISATION MATTERS TO USERS

It's no secret that personalised experiences have a greater ability than one-size-fits-all solutions to captivate users and support engagement. An interface that speaks directly to a user's preferences can command attention and prolonged interaction. In addition, platforms that understand and cater to a user's unique preferences feel less like tools and more like personal assistants, supporting users in achieving their goals – think of e-commerce recommendations or your social media feed.

MOVING TO HYPER-PERSONALISATION

Linear, one-size-fits-all experiences find themselves at odds with the essence of personalisation. Such pathways, fixed in their progression, often cannot flexibly accommodate the varied tastes and preferences of a diverse user base. The result? Users feel limited, rather than understood. Worse, users may feel left out across varying factors from accessibility to the lack of inclusive framing.
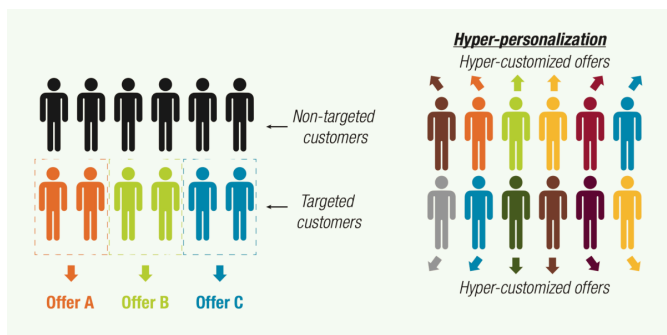
Source: Capgemini Consulting & ESSEC

INTRODUCING HYPER-PERSONALISATION

Many of our digital experiences are at the segmentation stage, with some – like popular social media or music services – reaching the micro-segmentation stage due to their impressive (both in positive and negative senses) proprietary content algorithms. Let's take the positive view for the sake of the discussion and ask ourselves, what's possible when we have enough data about a user to provide them the right content at the right moment? What's possible is for products and services to become more inclusive and for people who were previously excluded to be brought in from the cold. Building systems that flex to meet the needs of a diverse group of users, instead of building only core

experiences that apply to 80% of your users, allows brands to enhance their customer base and brand image. What this implies is a movement from micro-segmentation to delivering products that are unique to each individual. This is called hyper-personalisation and it is a tool that, when deftly applied (more on this later in the piece), allows products to become more inclusive and more accessible.

A HYPOTHETICAL EXAMPLE

Take a moment to imagine that we are in an alternate reality in which all pizza, whether frozen, fresh, professionally made, or handmade is sold (or licensed) by a company called Pizza Inc. In this world, Pizza Inc. is the most hyper-personalised organisation in the world for its ability to deliver a nearly infinite customisable pizza-eating experience. So if you live in Brisbane and go to your local restaurant to have a nice vegan tomato aubergine pizza for dinner or if you are in Lagos and you make pizza dough at home and prepare an ultra-thin crust margarita or if you are in Chicago and you get your favourite deep dish pizza delivered, you will be having the pizza experience that you desire. On the detail level, everything sounds different, but on the meta-level, they all shared in the pizza-eating experience and it was all delivered by Pizza Inc.

So how does Pizza Inc. deliver its service? Well, they serve a product that is effectively a canvas that can be customised in near limitless fashion as long as major rules are followed such as a dough base and toppings. Another key aspect of Pizza Inc.'s effectiveness is its massive amount of pizza-related customer data. They know your favourite pizza, they know the toppings you love, and they know the ones you hate. They also know your cultural context, how hungry you are, and your preferred pizza-eating method (home, delivery, restaurant) to deliver a pizza experience that fits your overall objective. With this information, Pizza Inc. delivers hyper-personalised pizza experiences around the world. Are you starting to get a bit worried about how Pizza Inc. deals with data privacy? Are you wondering if Pizza Inc. is trustworthy? Do you think having the perfect pizza experience is worth it? Let's unpack these questions in the following section.

ETHICS AND THE TRUST GAP

The concept of hyper-personalisation is scary for both companies, worried about liabilities, and customers

worried about privacy. How do we justify the need for it? Well, we consider real-life experiences such as getting advice from a teacher or guidance from a doctor. These are experiences that are rooted in a level of data-sharing but provide massive upsides and it all comes down to incentives.

INCENTIVISING TRUST BUILDING

Teachers, doctors, engineers, lawyers, and many more professions have both legal and internal rules and laws to incentivise individuals to act ethically and responsibly. They also have external social and economic factors that influence good behaviours. A lawyer who consistently breaches client confidentiality is not likely to become a financial success or be regarded as a good neighbour. Unfortunately, we have some negative externalities within the tech ecosystem that encourage bad ethical behaviour and we see that in customer data. According to Salesforce, only 51% of individuals trust companies. In the same report, we learn that 68% of customers say that AI's advancement makes it even more important that they trust the companies that use it. This means that delivering AI to customers must start with building user trust, maintaining transparency, and ensuring user safety before we consider experience, let alone profit.

A NOTE FOR FURTHER INVESTIGATION

There is more to consider on the topic of trust and ethics within the field of AI regarding model training and the working conditions of people who are essential to preparing training data sets as well as removing harmful content, but for this piece, we will focus on user experience.

SUMMARY

As the digital realm becomes increasingly saturated, the clarion call for designer practitioners and product experts is clear: embrace the multifaceted nature of hyper-personalisation. Genuine hyper-personalisation transcends mere cosmetic changes. It demands a symphony of content, tone, journeys, and layouts harmoniously orchestrated around a user's unique context and objectives. Understanding a user's unique context and objectives will undoubtedly require considerations about user data and will make trust-building the most important element of designing AI systems.

Ultimately, it requires a rethink from the segmentation approach of personas towards a new challenge, designing for an audience of one.

KEY TAKEAWAYS

Designers and product experts must adapt to the growing need for hyper-personalization in the digital world.

True hyper-personalization involves integrating content, tone, journeys, and layouts tailored to each user's specific context and goals.

Designing AI systems must focus on trust-building, considering the importance of understanding each user's unique context and objectives.

A shift from persona-based segmentation to designing for individual users is essential.

AI AS AN ENABLER

With the growing desire for hyper-personalised experiences, the question isn't formulated with 'if' these experiences will come to be, but 'how?' And, what challenges lay ahead? GenAI is poised to be the difference-maker here as a tool for infinite customisation and configuration and we, designers and product experts need to assume the role of a world builder. We need to focus our efforts on defining what's possible and what isn't rather than defining a set of defined actions to be carried out in a semi-linear fashion. In doing this, we need to embrace the weird and wacky (as compared to our well-researched approaches of the last 20 years) aspects of GenAI and access our humility over our egos in recognising that no one person knows what should come next, but this is where the most fun can be had, as we collectively attempt to define the near future.

A NEW PARADIGM

With the arrival of GenAI, we have seen a series of experiences that have changed our expectations for both what products and services do, but also what our role is in the task-completion process. Jakob Nielsen, the UX legend has posited that we are ushering

in the era of intent-based user interactions. This transition represents not just an evolution, but a revolution, in how humans and technology communicate. Instead of adhering to a machine's rigid language, users will express their desires, and the machine will orchestrate the 'how,' giving users a more intuitive and dominant role in the digital conversation.
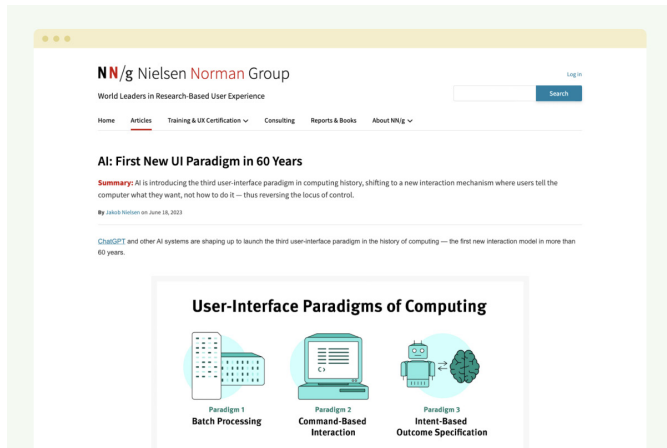
Source: Nielsen Norman Group

FROM ABSTRACTION TO INCLUSION

There is a sense that, as our digital ecosystems evolve, we as humans are becoming less so, that we are losing something, that we are wishing for times when we can use cmd-z in real life. This may hold some truth, but a prevailing fact is that GenAI is undoubtedly poised to make our digital ecosystems more human. This is because natural language is the new programming language, an abstraction that we have chipped away at since the ENIAC in 1945. Without practical need (our computers still communicate in 1s and 0s) for abstraction, new sets of individuals will be able to access and hopefully be helped through digital experiences.

WHY IS THIS BETTER?

It's tempting to wonder if we are at risk of reinventing the wheel when it comes to redesigning our digital interactions, but don't be fooled, if designers and product experts don't design the future, someone else will and we will all have weaker experiences because of it. So if we aren't reinventing the wheel, then what are we talking about? Why is this a better experience? Well, we are talking about a re-organisation of task definition and completion. In our previous paradigm we as the individual would be tasked with translating our intents into a series of steps to reach them. For example, 'I would like to cite my sources' in this

paper as an intent, yields a task list of 'select citation style', 'research the rules of the citation style', 'collect relevant information on referenced materials', which could be broken down into additional subtasks like 'find author', 'find publisher', etc. In our new paradigm, there is no need for this intent to be broken into tasks, rather as an individual your job is to define an intent and let the machine translate and execute tasks in whatever order using whatever logic it pleases and then you review the result. So you would say, cite the sources in this paper using Harvard style and the machine would produce a result and your job would be to inspect the work, requesting additional modifications to the intent as needed, such as 'remove the publishing location from all citations'. Knowing what Harvard Citation Style is (though it would be helpful) is no longer required to achieve your goal.
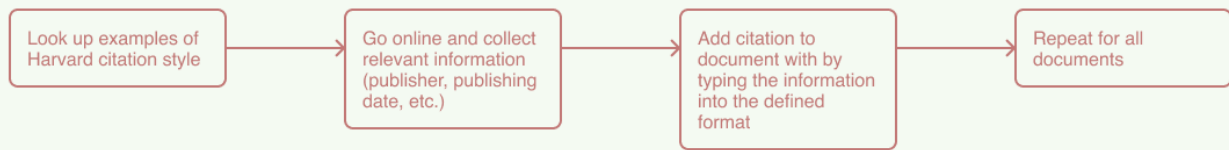
The reason this experience is better is because it is more human in its approach. You make requests the same way that you might to a colleague and your relationship with the work is the same as if a human added the citations and provided them for your review. This in turn will open up the often confusing digital world to many more users than ever before, whether they are asking their plants, how much water they need in a Plant-AI WhatsApp chat or desperately trying to configure their thermostat despite being without the manual en route to saving energy costs.

AI DOING ITS THING

Creating digital experiences that feel more human means that more humans will intuitively understand how to interact with them, but like real human interactions we have to tackle the challenge of dealing with background tasks or 'AI doing its thing.' At the time of this writing, even some tasks that feel basic to us, like describing a paragraph in one sentence can take quite some time, and if we asked a human to do this not only could they do it more quickly (in many cases), but they also provide a visual (and sometimes auditory) indication that they are thinking. They say, 'um..' to give themselves a moment before pressing on with their summary. Machines are not humans and they don't command the same expectation or understanding of the time needed to process. This is because doing a math problem on a handheld calculator somehow feels the same as asking Chat GPT to complete a task, 'it's just a machine' we think to ourselves. Add to the equation that we don't know what the AI might be doing while we wait, and you have a recipe for inpatient
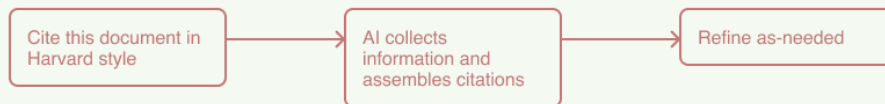
Command-Based
Add document Citation

| Look up examples of Harvard citation style | → | Go online and collect relevant information (publisher, publishing date, etc.) | → | Add citation to document with by typing the information into the defined format | → | Repeat for all documents |

Intent-Based
Add document Citation

| Cite this document in Harvard style | → | AI collects information and assembles citations | → | Refine as-needed |

and potentially confusion. Therefore, it is imperative that as we delve into designing these AI experiences we place great consideration on how we inform users what is happening in the background and how we let them know when outputs are ready for their review.
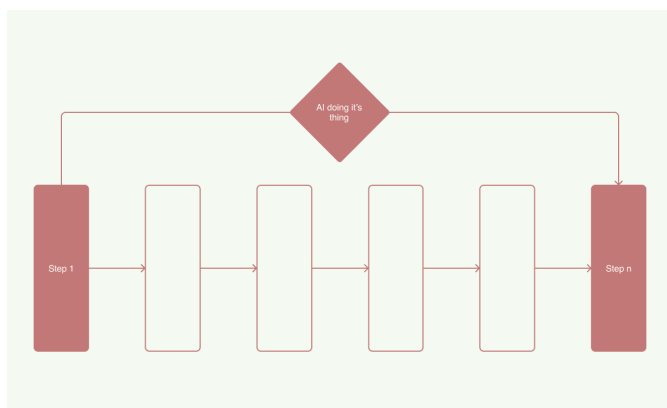

THE EXPERIMENTATION PHASE

When we look back at infinite scrolling or pull-to-refresh, they feel like wholly natural experiences with digital products, but we often forget that someone had to design those experiences and in that process, came many alternative solutions that were gradually phased out. So what is the dominant pattern of GenAI experiences? We simply don't know, but we can be certain, we can not get there without a healthy dose of experimentation.

WEIRD, THEN GOOD

Digital experiences have trended towards the aesthetic of slickness and an obsession with 'natural'

feeling interactions. Let's be kind to ourselves as designers and product experts and step back from the zeitgeist. We need to do this because it's time to suspend our understanding of the current moment and rather allow things to get weird before they get good. This means questioning every principle and allowing oneself to imagine new experiences without KPIs or your portfolio in the back of your mind.
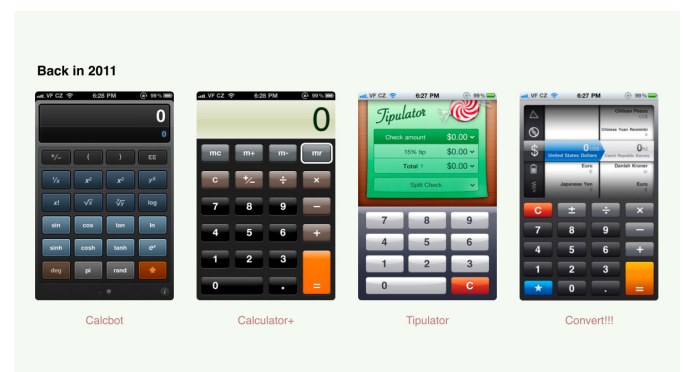
Source: Pttrns vis webarchive.org

To illustrate the moment we are in, we could look back to 2011 via the Way Back Machine and consider 4 of the popular mobile calculators of that time as presented by Pttrns. We notice immediately that there are some strong differences and designers are grappling with visual aspects like shadows, interaction aspects like multi-functionality, and meta-questions like realism. If you instead look just 2 years into the future you notice a newcomer has arrived. Completely flat with a swipeable drawer to select the currencies you desire to convert, it is by some margin the odd-calculator-out. This is what we need now, to define solutions outside of the dominant patterns of today so that we can investigate what can be possible tomorrow.

Currency – Simple Converter
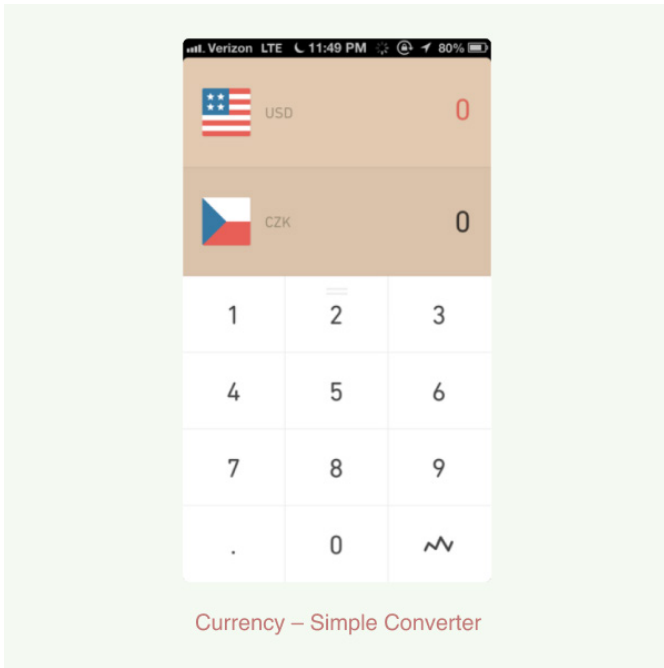
Source: Pttrns vis webarchive.org

THE PULL-TO-REFRESH MOMENT

David Haong summarises the situation clearly when he says, we are waiting for AI's pull-to-refresh moment – a defining user interaction that will set the standard for future interfaces. As we transition into the AI era, the challenge is to explore beyond conventional designs, awaiting that transformative idea. A whole world of new interactions and UX patterns is currently being discovered. What will be GenAI's 'pull-to-refresh' moment? We will only know through experimentation, collaboration, and conversation.

The next set of defining user interactions will need to touch on user interfaces and interaction design, but also on the relationship between the user and the machine. A storyteller might call this framing, and it will be essential to frame individuals' interactions with GenAI systems, whether as "individual and assistant" or "individual and co-pilot" or other combinations, to re-program individuals' expectations and set them up for positive experiences with our products.

FROM PLATFORMS TO OPEN WORLDS

It's hard to place this industry moment with a comparative opportunity, but one that springs to mind is the moment when video games moved from platformers (2D side-to-side movement) to open-world (3D freedom of movement) games. While much of the game logic and narratives are portable, there is a noticeable change in user agency. GenAI allows

us to make the same jump, but we must ensure we provide the right amount of constraints to keep users safe while encouraging them to enjoy the service or complete tasks using their desired approach.
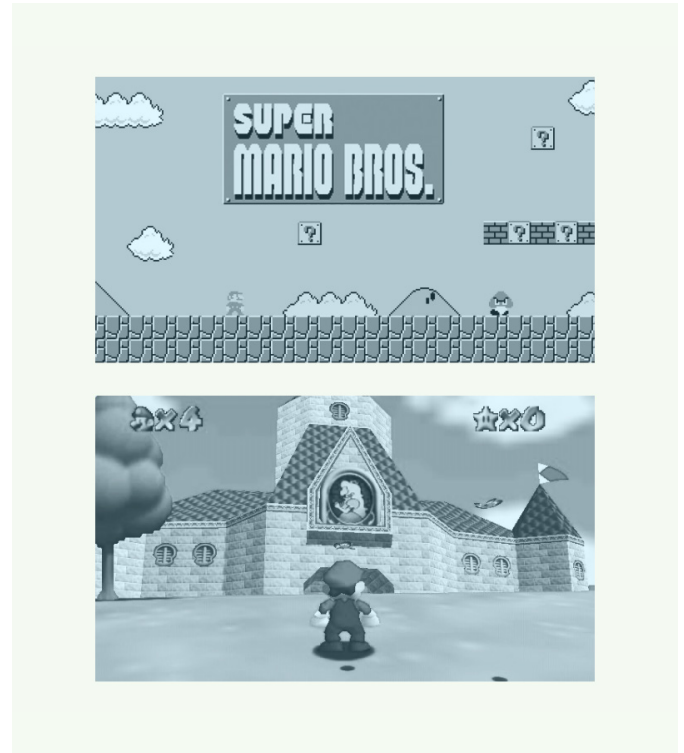
Source: Nintendo

CONSTRAINTS AS AN AGENT OF STORYTELLING

If you have ever set foot on a sports field, you will notice all sorts of strange markings on the grass. These are just one level of constraints to the players and they help shape the game, which helps narratives form and fan bases to unite around a team and a sport. We also see constraints outside of the world of sport, think of a director masterfully framing a shot for a film to fit a 16:9 screen or a fulfilment specialist changing an approach to packaging and helping get more products out of the factory in a more efficient manner. We as humans need constraints because they allow us to make sense of a situation and begin adapting.

As designers and product experts, we have a responsibility to build constraints into our experiences to give them context. Traditionally, we have done this by limiting the number of pathways available to individuals at any given time and we normally picture this as a nifty tree diagram. The challenge here is that interacting with an AI doesn't have an inherent path and each individual's objective and context will drive their flow. So we need to define certain rules of our system,

almost like walls in an open-world game, that define what's not possible, without ascribing a finite list of possibilities. This will allow the user to feel free without our system while maintaining their safety and our brand communication at all times.

RETURN TO PIZZA INC.

The pizza juggernaut is back and once again we require your imagination. Pizza Inc. exercises only 2 major constraints, all pizza must be on a dough base and all pizza must have a topping of some kind. With only these two constraints, the experience begins to have shape and individuals can start to make choices about what kind of pizza they would like. So as we consider what constraints to add, keep in mind, that it only takes a few to bring clarity and context to the experience.

OPEN WORLD UX

One way we might describe the systems we are discussing would be Open World UX. Open World UX has several distinguishable features, they are predictable and consistent, they afford clear narratives and they allow for tailored experiences. Open World UX experiences need to ensure that for each user input or action, there's a consistent and expected response, with variation (read accuracy) fine-tuned based on the scenario. It's like an artist who knows precisely which colour will emerge when two paints mix, but who in certain scenarios is open to nuanced surprises. Open World UX experiences also create clear narratives, ensuring that the user's journey has clarity, purpose, and direction. Finally, these experiences can be fine-tuned by designers by adjusting the constraint set, like a composer adding or removing instruments to the orchestra and therefore affecting what can be played and how.

For designers and product experts, these open-world experiences create opportunities for truly custom interactions, adaptive layouts, and immediate feedback loops. Designers can curate specific interactions for users, ensuring that each touchpoint feels personalised. Given the defined constraints, the design can dynamically adapt to user preferences. Finally, any changes or enhancements in design can be instantly reflected, allowing for real-time (non-predetermined) iterations based on user feedback.

SUMMARY

Open World UX experiences use constraints to bridge between the designer's intent and the user's aspirations. By dictating how different design elements come together, they facilitate an environment where experiences are both consistent and personalised, ensuring that each user feels seen, understood, and valued. As we continue to explore the nexus of AI and design, Open World UX experiences set a vision for a digital future that prioritises user agency while ensuring user trust and safety at all moments.

KEY TAKEAWAYS

GenAI heralds an era where users express desires, letting machines handle 'how'.

Clarity on background processes, such as notifying users how what's happening behind the curtain, mitigates confusion and impatience.

Embracing 'weird' experiences is pivotal in designing beyond our current design patterns.

## FROM LINEAR TO DYNAMIC JOURNEYS

The realm of UX has witnessed a myriad of evolutions, but one foundational approach consistently evident in many interfaces is the linear user journey. At its core, a linear user journey is a step-by-step progression, guiding users through a series of stages or events, generally towards a predefined goal. These journeys are often informed by comprehensive user research in combination with business goals derived from market conditions and organisational capabilities. The challenge is that organisations often miss out on opportunities to serve users at the edge of their customer base. This is because they believe it is not worth investing resources in services that touch only a small piece of the customer base. With GenAI in the role of orchestrator, designers and product experts have an opportunity to reach customers at the edges by offering dynamic journeys that morph to each individual's context and objectives. To embrace this we use modular approaches on our service and experience layers to ensure we provide users only the experience they desire, and nothing more, raising a further discussion on sustainable UX design.

LINEAR JOURNEYS

In the nascent stages of digital interfaces, many users found digital interactions confusing. Linear flows provide an easily understandable structure (for both user and machine), particularly beneficial for inherently sequential tasks, such as setting up a device or finalising an online purchase. The strengths lie in their simplicity, guidance, predictability, and translatability. Straightforward, clear paths eliminate confusion and help guide individuals to complete their goals. Additionally, linear journeys make it easy for individuals to predict what will happen next. Finally, these clear predictable steps can be easily translated into programmed functions, meaning we can build them.

CONSTRAINTS OF LINEAR JOURNEYS

True personalisation, a burgeoning necessity in modern UX, grapples with the constraints of linear designs. A linear journey inherently struggles to cater to individualised preferences, behaviours, and contexts. For instance, while a music app with a linear approach might guide a user through a series of steps to generate a playlist (create, name, add tracks, etc.), a dynamic model would be only one step (create playlist) that takes into account user preferences as well as external factors, such as trends.
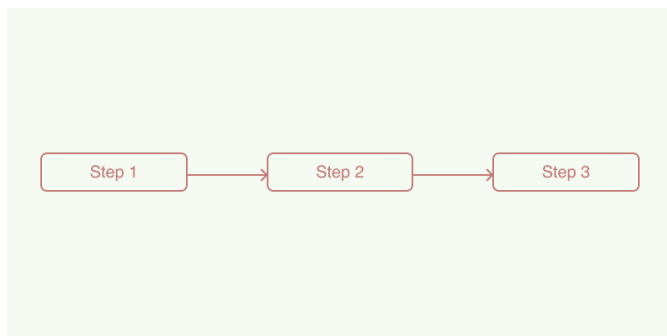

MOVING TO THE META LEVEL

While linear journeys in UX played an indispensable role in digital design's early days, the contemporary digital environment demands a more adaptive and individualised approach. We should see linear journeys as the meta-level of user interaction with our product or service. For example, if we have a linear journey that defines the major steps of a user purchasing a new coat (browse, select item, purchase), within those meta steps we can consider what can be dynamic (each browsing in their preferred way, such as by voice, us-

ing recommendations, leaning on trends, etc.). Our movement towards dynamic journeys, powered by AI, promises a 'humanisation' of the 'user', no longer aggregated into personas or target groups, but free to shape their own experiences.

DYNAMIC JOURNEYS

Dynamic user journeys are key to achieving hyper-personalised experiences. These journeys are coordinated by an AI that uses an individual's context and objective to assemble a productive experience on the fly. These journeys are hyper-personalised, objective-driven, and ephemeral. When considering how this works, familiar questions arise around AI Doing Its Thing and the simplicity and predictability achieved by linear journeys. Here, we must achieve the appropriate balance of dynamism, wrapping our dynamic experiences inside meta-linear journeys.
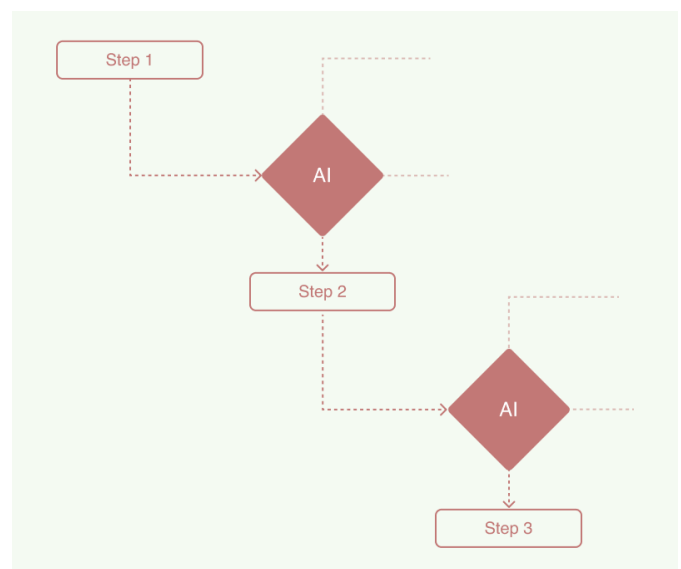
AN EXAMPLE

Imagine you would like to go to your nearest train station and you have only a paper map at your disposal. You would likely locate yourself and then consider what roads, or transport systems might take you to the train station. You may even draw your route on the map. Additionally, you would probably pay detailed attention to relevant street names so that you can orient as you go. In the end, you would have a set of directions that will take you to the train station and then you would set off and follow your plan.

In contrast, what if you had a GPS? In that case, you would enter your destination of 'Train Sta-

tion' and then review and select your travel options (public transport, walking, biking, etc.). Then you hit start. If the unexpected happens such as traffic or train delays the GPS immediately recalibrates in real time. Despite unexpected changes your focus, rather than on street names or subway stops, is on your destination. This is a dynamic journey.

OBJECTIVE-DRIVEN EXPERIENCES

Our dynamic user journeys are directly aligned with Nielsen's assertion that we are in the era of Intent-Based Outcome Specification. Borrowing from game design, we can consider each user's journey as a set of objectives or 'levels' to achieve and our linear meta journeys as the overarching narrative. As individuals progress, the system adapts, providing them with the right tools, content, and pathways. It's akin to progressing through a game where each level offers a unique set of challenges tailored to the player's skill.

This implies a new challenge for designers and product experts, supporting users to define their objectives while contextualising background processes undertaken by the AI system. Understanding objectives allows the product to use only the necessary resources. Additionally, energy conservation can be accounted for, delivering different content depending on the network type, energy mix and cost. This will allow for sustainability to be considered in the design of digital experiences. As Tim Frick explains in Designing for Sustainability, "Products and services that provide a streamlined yet enjoyable experience - putting the

right things in front of users at precisely the moment needed and nothing more - are more efficient and greener." While there are many questions regarding AI's environmental impact it's worthwhile to notice areas where it has an upside.

A MODULAR APPROACH

In a digital age where every user seeks a personalised experience, relying on fixed templates is like attempting to fit unique puzzle pieces into one generic mould. While templates offer a structured foundation, they have the potential to fail to capture the nuance and individuality users crave. Modular design emphasizes compartmentalized functionality. Each distinct module addresses a specific role, and when synchronized, these individual parts merge into a complex, adaptive structure, perfectly poised for modern digital experiences.

MODULARITY AND DYNAMIC JOURNEYS

What does a generative design system look like? How can components adapt depending on what's nearby or when a user's context changes? These are the questions that await interaction designers over the next years. A modular infrastructure simplifies our ability to be adaptable, which is essential to delivering dynamic hyper-personalised experiences.

Modular design, amplified by AI's prowess, is reshaping our digital landscapes from static templates to fluid, constraint-based, user-centric canvases that adapt and evolve with each user interaction. This promises not just a responsive interface, but a revolutionary one, where users aren't just participants but co-creators of their digital narratives.

SUMMARY

Open World UX has helped us rethink how the service layer of digital experiences can be reimagined and now dynamic user journeys complete the puzzle by showcasing AI's potential on the experience layer. Linear journeys are not retreating, but rather become key in defining the meta experience that users should have, while dynamic journeys, placed inside larger meta journeys, define the moment-to-moment interaction with the system. Finally, to facili-

tate this on the user interface level we need to look to modular design and consider how we can allow greater flexibility within our components en route to creating truly hyper-personalised experiences.

KEY TAKEAWAYS

Linear journeys, now move to the meta-level and give way to dynamic experiences in which the AI plays the role of conductor, adapting user journeys to individual contexts and objectives.

Linear journeys, while essential, struggle with true personalization, lacking adaptability to varied preferences, behaviours, and contexts.

Dynamic journeys, akin to using a GPS, react to real-time changes, focusing on objectives over predefined steps.

Modular design, bolstered by AI capabilities, enables fluid, adaptive experiences that empower users to shape their digital narratives.

CASE STUDIES

Moving from theory to practice we will look at two case studies that illustrate the challenges and considerations when designing AI experiences. First, we will look at ProTunes One, a music licensing platform for content creators, which uses an AI-driven search to help users sort through their library of 2 million+ tracks. We will follow this up with a discussion of NUI Care, a mobile application for caregivers in Germany who care for their older loved ones. Nui's AI assistant helps translate complex bureaucratic care benefits information to provide users with concise and clear information about what benefits they can receive as well as other key information to improve their care routine.

PROTUNES ONE

Protunes One is an AI-driven B2C solution specifically designed for content creators in the field of music production. The application focuses on streamlining music licensing for platforms such as YouTube and other social media. We will look at two aspects of our work on ProTunes, prototyping and managing user expectations.

PROJECT CHALLENGE

How might we help users find music that is unique to them – faster?
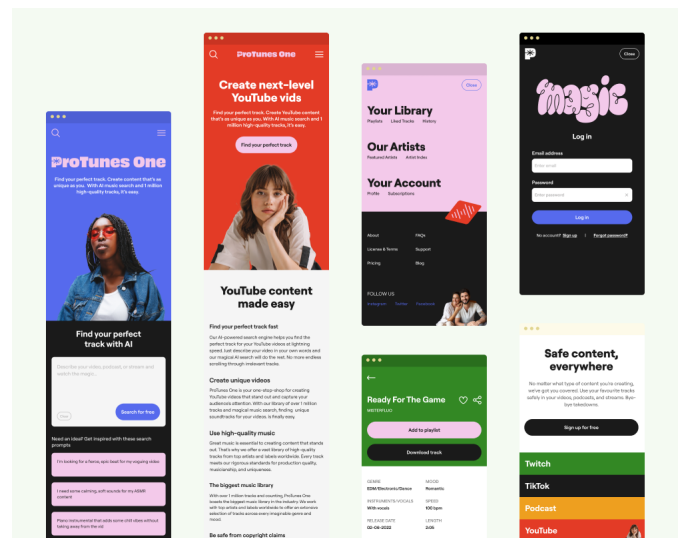
PROTOTYPING

Prototyping, especially when developing new patterns and experiences is essential to building quality digital products, but when prototyping AI experiences things get tricky. As designers we are used to providing static clickable prototypes and running user tests, but how do we do this with AI systems? It's not enough to test only the interaction design, we must also interrogate the performance of the service.
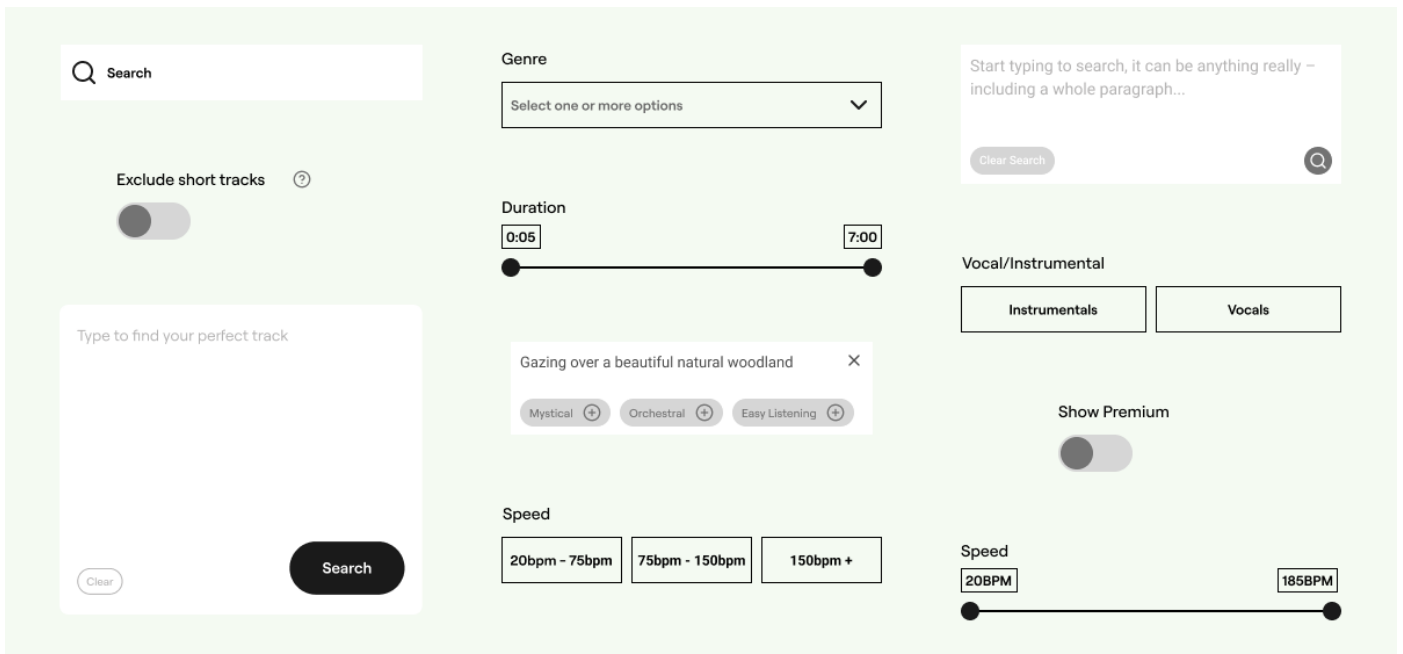
To do this on Protunes we developed two prototypes, one static clickable and one technical. In our static prototype, we focused on information architecture, affordances, and interactions. In our technical prototype, we focused on assessing the performance of our experience (i.e. how well the songs returned from the search match the user's expectations).

TAKEAWAY

Prototype in parallel to assess the model and the interaction to leverage both the current skills of product teams. Additionally, we need to evaluate what tools we need in the future to deliver valuable research results for these experiences.

MANAGING USER EXPECTATIONS

A second challenge in designing AI experiences is managing user expectations, which are changing rap-

idly as new AI experiences are deployed across industries. In the case of ProTunes, we worked to understand how we could make users comfortable with our AI music search experience. This meant providing users with common search parameters such as genre and mood, but also trying to educate them on how the free-text AI search could provide more accurate results more quickly.

What was most challenging was how to resolve the tension points between search results from metadata (i.e. this is a 'hip-hop' song) and results from the AI search which identifies hip-hop tracks based on their musical qualities. This situation accurately reflects the reality that we are in a transitional moment between our current approaches and the new opportunities afforded by AI. We did not explicitly 'solve' this tension, but rather we allowed both approaches to track sorting to work separately or together.
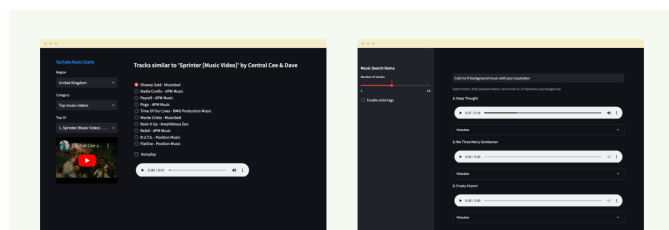

EXAMPLE

Why might we need a manual override such as a filter, if our claim of a quick efficient and accurate AI search is true? There are two reasons to provide this that are best illustrated in an example. Riding the train in Germany (and in many other places) you may notice that the doors between cars are automatic. If these doors

are working you would not notice them, but the moment they malfunction they suddenly sit in the forefront of your mind. Luckily, there's not much to fear because, in the case of the Deutsche Bahn, each door has a visible sensor that a person can activate with a wave of their hand to get the door to open. How do we even recognise this sensor with its shiny black face and single red eye? Well, we know them from experience, from airport toilets to sinks in restaurants. With this collective knowledge, a door is relatively easy to 'debug' and move on with your life. This element of debugging is not yet explicitly designed into our experiences with AI systems and will need time to design them and for the solutions to seek into our collective understanding.

TAKEAWAY

Provide outlets for users to manually control the system when things don't go as planned.
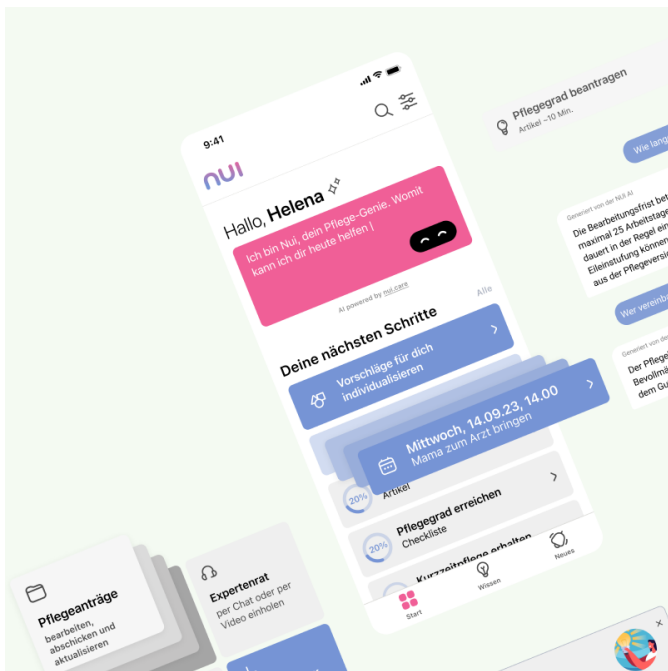
NUI CARE

Nui Care is on a mission to support caregivers and in turn, support the ones they care for. This current project emphasises the crucial role of Generative AI as a tool for creating personalised care support. In this case, we will consider team setups, ideation and the role of modularity in delivering a high-quality care support experience.

PROJECT CHALLENGE

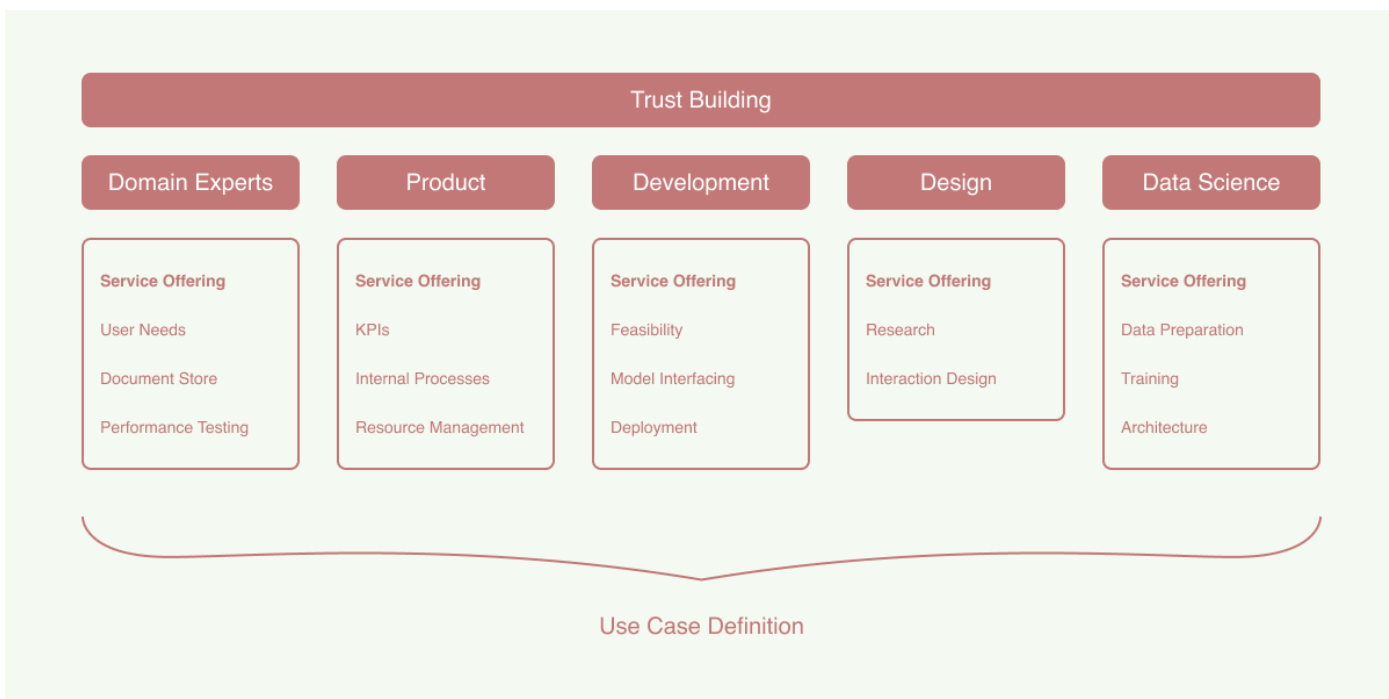How might we provide care expertise at scale and ease the burden of caregivers?

BUILDING THE RIGHT TEAM

When bringing Nui from idea to proof of concept to product we needed to have the right team. In our case to build an effective product we brought together data scientists, designers, developers, product managers, and domain experts. While most digital products function with a core team across product, design, and development, GenAI projects require tight-knit collaboration with data scientists and domain experts to make the initial idea a reality.

DE-RISKING THROUGH COLLABORATION

One key learning from Nui was recognising that it is no longer possible for a designer to completely orchestrate and design an experience. This is because the AI layer has such a prominent role the the user's experience. No number of intuitive interactions can save a poorly trained AI assistant. This makes it all the more important that each team, product, design, development, data scientist, and domain expert works together to de-risk the work of one another.

In Nui's case, we were eager to provide easily understandable explanations to complex care topics – an excellent use case for a large language model. One challenge we faced was, how to express a level of credibility to our users. How do we make them feel comfortable with the assistant communication? and finally, how do we support users when they are unsatisfied with the answer coming back from the assistant?

Trust Building				
Domain Experts	Product	Development	Design	Data Science
Service Offering User Needs Document Store Performance Testing	**Service Offering** KPIs Internal Processes Resource Management	**Service Offering** Feasibility Model Interfacing Deployment	**Service Offering** Research Interaction Design	**Service Offering** Data Preparation Training Architecture

Use Case Definition

Every team has a role to play in answering these questions and depending on your resources you may opt to deploy varying solutions. In Nui's case, we could ask the data science team to undertake expensive training of our assistant, continuously fine-tuning the performance until we reach some agreed-upon level of quality. We could also ask our domain experts to provide more detailed documents about the various care topics users are asking about. We could ask the product to rearrange resources to allow more to flow to data science for the training. We could ask the development team to provide useful and actionable widgets for each answer. Or we could take the approach we used, which was for the design team and the data science team to discuss and agree to provide citations to the users so they could understand which document Nui's assistant was using to generate the answer. This showcases the role of recognising that every team has an essential role to play in delivering high-quality AI experiences.

TAKEAWAY

Diverse teams with equal power dynamics across disciplines de-risk GenAI use cases and make it more likely that you will bring the right product to market.
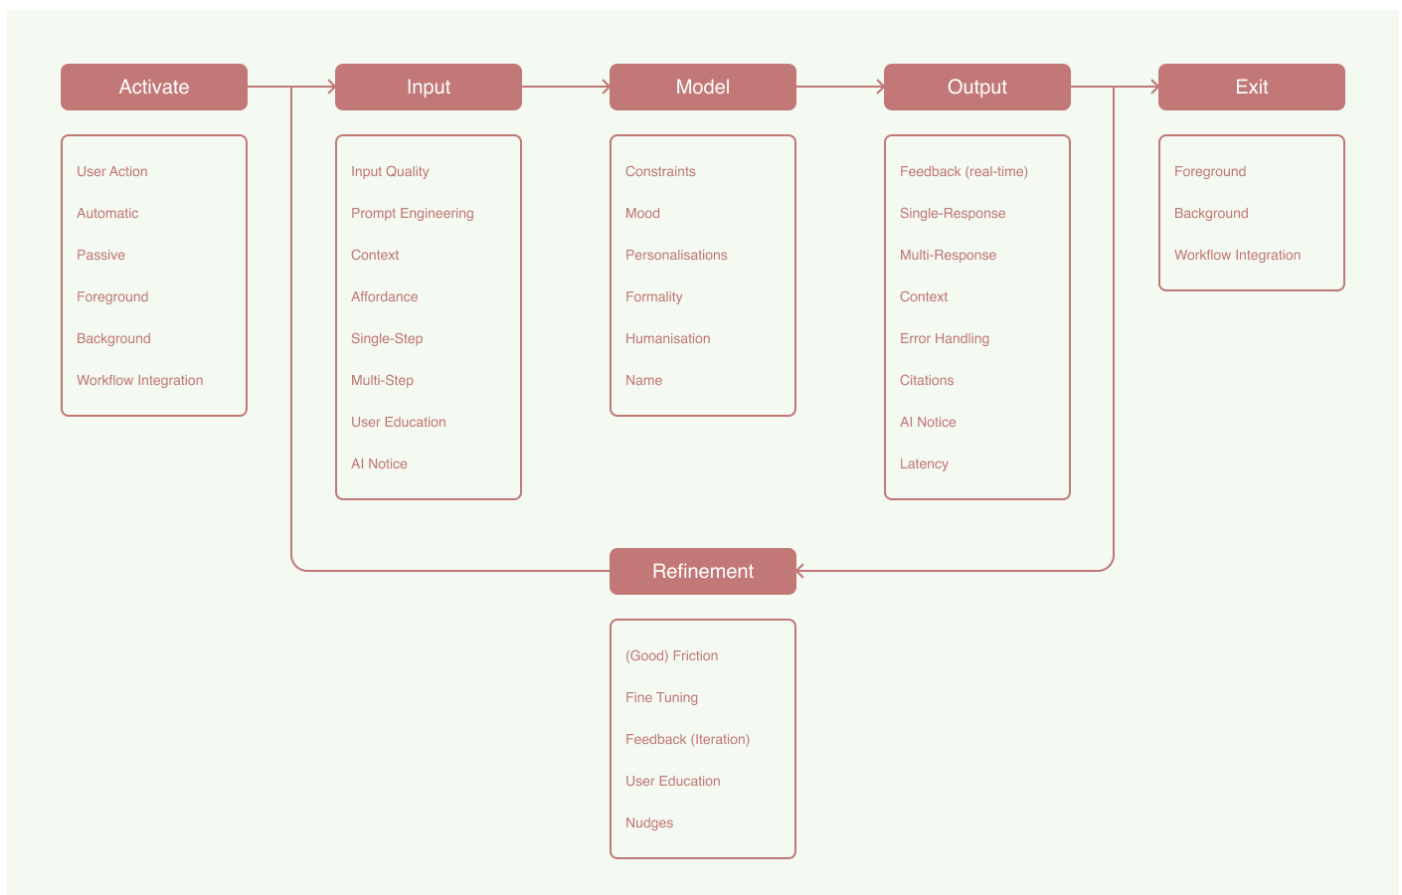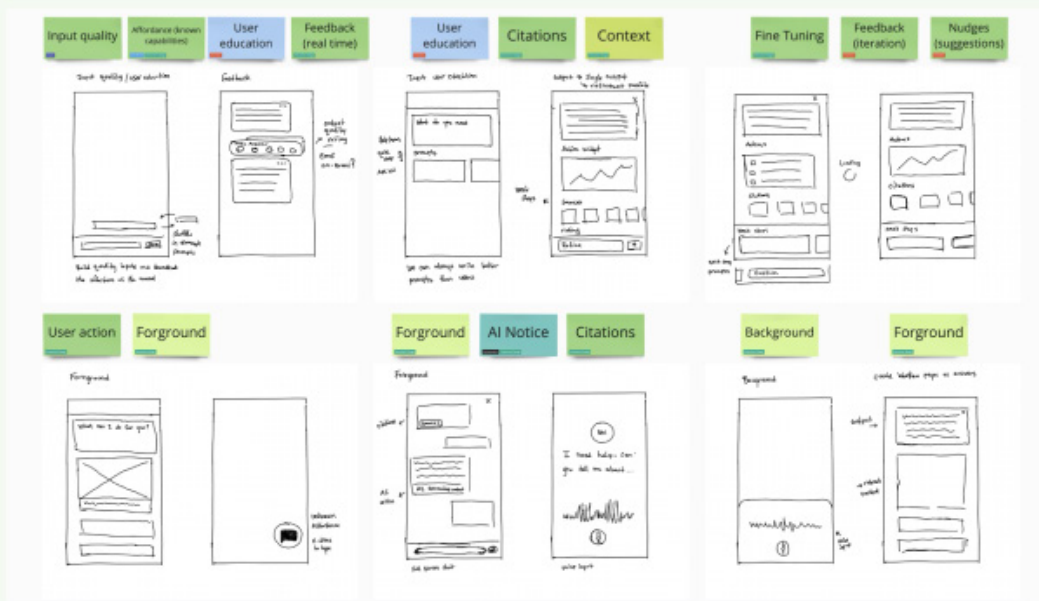
IDEATION

Once the right team is in place the task turns to coming up with a set of strong ideas to prototype, test, refine, and build. We know that there are new AI use cases every day from corporations and enthusiasts alike. To prevent us from feeling overwhelmed and to centre our focus we developed an ideation framework that we used as a tool for building up small ideas that we could eventually remix into concepts.

THE FRAMEWORK

To frame our ideation we first defined an overarching goal, and that's building user trust. Below this, we define a cycle of a user's interaction with an AI system. This moves from Activate to Input, to Model, to Output, to Exit. After the Output is provided there is also a possibility of moving to Refinement. Below each moment in our cycle, we have sets of design qualities and areas to pay attention to in a given moment.

To use the framework we selected a state such as Activate and then we selected 1 - 3 design qualities. With those in hand, we asked ourselves questions like, "How could we design an AI-driven care experience that can

be activated passively and with a high focus on work-flow integration?" Then we sketched ideas. We continued to remix qualities and look at different areas of the cycle until we had many ways to provide what we believed was a quality experience. Then we remixed our smaller ideas into larger concepts en route to developing prototypes ready for user testing.

SUMMARY

Each of these cases showcases the nuances of designing AI experiences and makes clear that there is much work to be done by designers from across regions and industries to help make sense of what the future will look like. Whether it's new collaboration models, different approaches to ideation or deep thoughts on the tension between the 'old' and the 'new' we as a design and product community need to be having these conversations, experimenting, and sharing our learnings.

REFLECTION AND CALL TO ACTION

The field of UX has seen major changes over the years. In this paper, we explore the move from pre-determined user journeys to personalized, AI-driven experiences, highlighting the new challenges and opportunities for businesses and designers. Ahead we consider what worries us and reiterate our reasons to be optimistic.

CRITIQUES IN IMPLEMENTING HYPER-PERSONALISED EXPERIENCES

While the opportunities ahead should provide a jolt of inspiration for design practitioners and businesses, it must also be acknowledged that genuine challenges remain to obtaining this desired future. Concerns exist across high-level topics like user privacy and the role of computationally intensive computing on climate change to the practical details of how we should test and monitor services, as well as how to educate users on how the system works. This is about balancing abstraction, giving users distance from technical mechanisms so that they can focus on their objectives, and transparency, ensuring that users understand what's going on beneath the surface and why.

Here's a breakdown of the key short-term challenges:

Testing: The dynamic nature of rules-based experiences can make conventional testing methodologies seem obsolete. Testing such experiences requires not only validating the individual modules but also understanding how they dynamically interplay with various user inputs and contexts. This adds layers of complexity in ensuring that the experience is consistent and error-free across multiple user journeys.

Monitoring and User Data Privacy: Constant monitoring is crucial for these systems to learn and adapt. However, this brings forth the challenge of how often

and how deeply a system should monitor user interactions. Balancing this necessity for improvement with respecting user privacy and adhering to data protection regulations becomes a tightrope walk.

User Education and Transparency: As we move away from linear experiences, there's an increased onus upon businesses to educate their users. They need to be made aware of how the AI model operates, its continuous learning mechanisms, and the implications of this learning on their future interactions. This isn't just about trust but also about enabling users to get the best out of the system.

Conventions and Brand Consistency: The challenge here is twofold. On one end, businesses must ensure that despite the adaptable layouts and dynamic content presentations, users can still recognise and relate to the brand. On the other end, there's a need to establish conventions that ensure users, even when presented with novel interfaces or content layouts, can predictably navigate through them.
Impact on Climate Change due to Computationally Intensive Computing: The push towards more advanced AI and real-time personalization also means more computations. These computationally intensive operations, especially when scaled globally, have a significant carbon footprint. With growing concerns about climate change, there's an urgent need to develop eco-friendly computational solutions or offset the environmental impacts in some meaningful way.

Balancing Personalization with Generalization: While the drive is towards hyper-personalization, there's a risk of making the user experience too niche. This could alienate potential new users or limit the breadth of content/experiences available to current users. The challenge is to strike a balance, ensuring the system can cater to both broad and niche user segments effectively.

Navigating the intricate maze of these challenges necessitates a blend of technical acumen, ethical decision-making, and a relentless focus on user needs and expectations. The journey promises transformation but is lined with intricacies that require careful navigation.

CALL TO ACTION

The digital era began with static, linear pathways, a one-size-fits-all model which, though functional, didn't account for the nuanced needs of diverse user bases. This static nature often felt limiting, curtailing the potential for immersive, engaging experiences tailored to each user.

Our exploration revealed the need for a shift—a move from these restrictive pathways to a landscape defined by adaptability, precision, and personalization. We've demonstrated the immense potential generative AI holds in reshaping this landscape. Systems like ChatGPT and others herald a new dawn in UX, one where intent and outcome are prioritized over procedural command, mirroring Jakob Nielsen's idea of an "intent-based outcome specification."

However, we argue that the future lies not just in leveraging AI capabilities but in marrying them with robust, constraint-based frameworks. Such frameworks, informed by user needs and iteratively refined by feedback, can generate experiences that feel both tailored and organic. This shift is monumental, as it introduces a UX where users no longer feel like they're navigating pre-defined paths but are co-creators in their own digital journeys.

Throughout this exploration, our discourse has been deeply rooted in hands-on experience, ensuring that our insights and propositions are not merely theoretical extrapolations but are borne out of genuine challenges and solutions from our active involvement in a caregiving project. The melding of GenAI with UX, as showcased in our cases, is not a distant future but an ongoing revolution.

Yet, as with any transformative shift, challenges lie ahead. AI's vast potential is counterbalanced by issues of transparency, control, and ethical considerations, particularly in how user data is managed and utilized. As we pivot towards this promising horizon, businesses, developers, and designers must ensure that the experiences crafted prioritize user trust, safety, and understanding.

So, we conclude with a clarion call: to industry professionals, to UX pioneers, and to businesses. Embrace this shift, invest in the confluence of GenAI and UX, and harness its potential responsibly. The result? A future where digital experiences aren't just interfaces, but where they are adaptive, intuitive extensions of human intent, bridging the divide between technology and the individuals it serves.

ACKNOWLEDGEMENTS