MASTERARBEIT | MASTER'S THESIS

Titel | Title

Blind Spots: Deceptive Patterns and User Perception in Commercial Map-Based Search

verfasst von | submitted by

Laura Welz BEd

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Master of Education (MEd)

Wien | Vienna, 2026

Studienkennzahl lt. Studienblatt | Degree programme code as it appears on the student record sheet:	UA 199 507 510 02
Studienrichtung lt. Studienblatt | Degree programme as it appears on the student record sheet:	Masterstudium Lehramt Sek (AB) Unterrichtsfach Englisch Unterrichtsfach Geographie und wirtschaftliche Bildung
Betreut von | Supervisor:	Univ.-Prof. Dr. Krzysztof Janowicz

Abstract

Digital Maps have become a paramount navigational tool of everyday life. Yet as much as we rely on them, rarely do we question if everything that is shown on maps is a factual representation of reality. Increasingly, maps have also become an influential actor within commercial platforms, as companies like Booking.com embed interactive map interfaces as pillars of their search and recommendation systems. This thesis examines how map-based search is transforming under conditions of commercialization and algorithmic ranking, and how it affects user perception, trust and decision-making. The study combines a cross-platform interface audit with a user survey to identify and compare deceptive or manipulative design patterns in map and list views across selected commercial platforms. Building on existing deceptive pattern literature, it first develops a map-specific typology tailored to spatial interfaces, which then serves as the analytical framework for the empirical investigation. The findings show that users' expectations of spatial completeness align only with what is visually explicit, while 'hidden' mechanisms such as preselected results largely remain unnoticed. Across platforms, deceptive patterns operate primarily through selective visibility rather than overt pressure, creating a strong illusion of user control. At the same time, the growing shift toward ranked lists and AI-driven recommendations further centralizes decision-making and reduces transparency in spatial search. By conceptualizing and empirically grounding a map-specific typology of deceptive patterns, this thesis contributes a framework for critically assessing commercial map interfaces in the AI era.

Kurzfassung

Digitale Karten sind zu einem unverzichtbaren Navigationsinstrument des Alltags geworden. So sehr wir uns jedoch auf sie verlassen, hinterfragen wir selten, ob alles, was auf Karten dargestellt wird, einer faktischen Abbildung der Realität entspricht. Zunehmend haben sich Karten zu einflussreichen Akteuren innerhalb kommerzieller Plattformen entwickelt, da Unternehmen wie Booking.com interaktive Onlinekarten als zentrale Bestandteile ihrer Such- und Empfehlungssysteme integrieren. Diese Arbeit untersucht, wie sich kartenbasierte Suchprozesse im Zuge von Kommerzialisierung und algorithmischem Ranking verändert haben – und welche Folgen das für Wahrnehmung, Vertrauen und Entscheidungsfindung der Nutzenden hat. Die Studie verbindet ein plattformübergreifendes Interface-Audit mit einer Nutzer:innenbefragung, um täuschende oder manipulative Designmuster in Karten- und Listenansichten ausgewählter kommerzieller Plattformen zu identifizieren und gegenüberzustellen. Ausgehend von der bestehenden Literatur zu „Dark Patterns“ wurde zunächst eine kartenspezifische Typologie entwickelt, die auf räumliche Interfaces zugeschnitten wurde und als analytischer Rahmen für die empirische Untersuchung diente. Die Ergebnisse zeigen, dass die Erwartungen von Nutzenden an räumliche Vollständigkeit nur mit dem visuell Offensichtlichen übereinstimmen, während „verborgene" Mechanismen – etwa vorausgewählte Suchergebnisse – weitgehend unbemerkt bleiben. Plattformübergreifend entfalten manipulative Designstrategien ihre Wirkung vor allem durch selektive Sichtbarkeit statt durch offenen Druck, wodurch eine starke Illusion von Kontrolle entsteht. Zugleich zentralisiert die zunehmende Dominanz von Listen und KI-gestützten Empfehlungen die Entscheidungsmacht von Nutzer:innen und verringert die Transparenz räumlicher Suchprozesse. Durch die konzeptionelle Entwicklung und empirische Fundierung einer kartenspezifischen Typologie täuschender Designmuster leistet diese Arbeit einen Beitrag zur kritischen Auseinandersetzung mit kommerziellen Onlinekarten im Zeitalter der künstlichen Intelligenz.

Acknowledgements

I would like to thank my supervisor, Univ.-Prof. Dr. Krzysztof Janowicz, for his guidance and insights throughout this project. I would also like to thank my family and friends for their encouragement and support throughout my studies.

Table of Contents

List of Figures

List of Tables

1. Introduction

Digital maps have become an indispensable part of everyday life. Whether booking a hotel, finding a restaurant, or checking public transport, it is almost impossible to interact with a digital service without encountering a map. In the past decades, standard online maps have transformed from simple navigational aids into hubs of information and commercial ecosystems. This shift reflects broader technological advancements, developments in user experience (UX) design, and increasingly complex algorithmic logic; in the case of machine learning models in particular, even their creators cannot always explain exactly how outputs are reached (Burrell 2016: 3). Many modern web maps are also deeply entangled with the commercialization of spatial data and monetization strategies. As platforms like Google Maps, Apple Maps & Co. guide, filter, and predict user behaviour, results are increasingly steered to show sponsored content or be more personalized. This is not necessarily a bad thing; however, it does raise questions about ethics, transparency, and user autonomy. In connection with the development of interface design, persuasive strategies have become a central aspect. In many cases, these strategies cross into manipulative territory. Brignull (2010) first coined them as *dark patterns*, or *deceptive patterns*, which describe interface decisions that pressure users into making certain decisions, usually for commercial gain. Examples include hiding additional fees until checkout, emphasizing urgency with fake scarcity claims, or preselecting certain options.

Research Gap

While these tactics have been well-documented in e-commerce, their presence in digital maps has received little to no academic attention, despite maps being particularly well-suited to such manipulations, given that maps, by their very nature, have always been tools of influence and control (Harley 1989: 2). Existing research on dark patterns has largely concentrated on conventional e-commerce environments and not considered the spatial dimension of interface. Moreover, little is known about how users perceive such patterns in map-based interfaces, which design choices they expect or accept, or how awareness diverges from actual interface practices. Recent developments further indicate a shift from maps toward ranked lists and recommendation feeds, some of which are already AI-driven, yet the implications of this transition remain underexplored.

Research Questions

This thesis investigates how digital map interfaces may facilitate or conceal deceptive design patterns, particularly as maps are increasingly influenced by commercial interests. It examines both the presence and visibility of such patterns, as well as the expectations users bring to map-based search experiences. Additionally, it also explores the growing phenomenon of maps being replaced by lists, which, in some cases, are already driven by AI. The exact research questions this study addresses are:

1. To what extent do user expectations align with the presence of deceptive patterns in map-based web interfaces?
2. Which potentially deceptive design patterns are commonly used in commercial map interfaces? Which of those patterns are users aware of?
3. How does the shift from spatial maps to (AI-driven) recommendation lists affect user control and transparency in location-based search?

Outline of Approach

These questions respond to the current changes in online maps. Most users may believe certain design interface choices to be “normal” and not realize that the visibility of listings, pins, filters, etc. serve business interests as well. For this reason, this thesis hopes to make three core contributions: First, it proposes a map-specific typology of deceptive patterns, adapting familiar e-commerce patterns to spatial interfaces and clarifying when they occur in map vs. list views (e.g., labelling of sponsored pins, default sorting etc.). Second, a user survey was conducted that captures expectations and trust signals in map-based search. It aimed to identify which interface choices people support, which they notice, and where opinions diverge. Third, it presents a cross-platform analysis across three categories: Hotels & Accommodations, POI & Reviews, and Real Estate; using a transparent 0/1 coding scheme (with brief evidence notes) to compare the visibility and prevalence of patterns across views and categories. Taken together, the findings point to practical steps for design, policy and research. The following chapters begin with a literature review, outline the methodology, present the results, and discuss implications and limitations. Importantly, this thesis does not seek to establish intent behind the identified patterns; whether they arise from purposeful design decisions or from technical and commercial constraints falls outside the scope of this study.

2. Literature Review

2.1. The Evolution of Digital Maps

It is in the nature of maps to always evolve and transform, but for a long time, cartographers focused mostly on charting the unknown, to create a complete picture of the world. This has changed over the past decades, with rapid advances in satellite imagery and surveying techniques, the Earth's surface has now been extensively mapped – although this process is still ongoing, with continuous refinements in resolution and detail, particularly in remote or rapidly changing environments. However, with this development also came a shift to *thematic* maps. Even though they have existed since the early 16th century, it was not until the widespread availability of human demographic data, that thematic maps truly started to take off (Muehlenhaus 2014: 4). Tied into this is the emergence of the modern nation-state. Through compulsory censuses, economic trade records, and administrative auditing bodies, sufficient data now became available for large-scale mapping (Muehlenhaus 2014: 4). Then, with the rise of accessible and high-speed internet, many maps became digitalized. Data, files and programmes were stored online, computing was done online, and online mapping became popular. During the late 1990s and early 2000s, Geographic Information Systems evolved rapidly and started expanding beyond specialized government and academic use, becoming widely adopted by businesses, municipalities, and the general public (Ali 2020: 4). In 2005, Google Maps was introduced to the world and changed the way we create and think about maps. Apart from finding directions, suddenly everybody could create maps by adding their own data layers and placing it over Google's base maps. Additionally, tools such as the panning and zooming, clicking on pictures and opening links were introduced (Peterson 2014: 409), thus turning online maps into *interactive* maps.

An important factor in this development is the graphical user interface (GUI), which determines how users interact with maps. While different mapping services across the world (e.g., Google Maps, Waze, Baidu Maps) have their own set of button graphics and arrangements (Medynska-Gulij et al. 2021: 2), a sort of standardized interface has emerged. Zejdlik & Vozenilek (2025) conducted a comparative study of eight popular web mapping services, evaluating their design, usability, and adherence to cartographic principles. Their findings reveal that while each application possesses unique characteristics, they share common interface features, such as tile-based zooming and panning. However, differences in map layout, labelling, and analytical tools can greatly impact user experience and interpretation. Beyond standardization, Medyńska-Gulij et al. (2020) examined how different interactive cartographic techniques influence perception, particularly in dynamic environments such as large events. Their study suggested that reducing certain UI (User Interface) elements – like scale bars and north arrows – can, in some cases, enhance clarity, particularly when they do not contribute to the interpretation of movement patterns. However, their findings also highlight that the effectiveness of UI simplifications depends on the specific mapping context, as experts held divergent opinions on which visual components were essential for analysis (2020: 13). This is indicative of a broader trend in digital mapping, where platforms are increasingly integrating real-time updates and adaptive map designs to optimize user experience. But while standardized interfaces are convenient for ease of use and familiarity, it also means that people have gotten used to a certain style of map and have started expecting all other maps to follow the dominant theme. "Any map that does not meet this expectation will not be acceptable. Conformity can both be a strength and a weakness, even Google continually modifies the design of its maps, but only slightly." (Peterson 2014: 409). For example, the same mapping platform will look visually different depending on the region it displays, as maps inevitably reflect the geographic and urban reality beneath them. This variation can be observed in aspects such as street density, transport networks, and labelling conventions. *Figure 1*, of Tokyo, and *Figure 2*, of Phoenix, Arizona provide a good example.

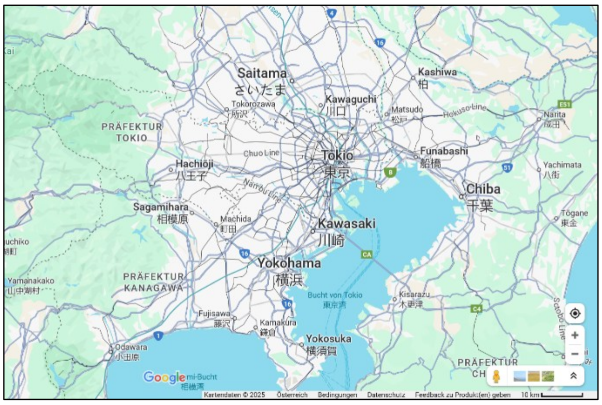

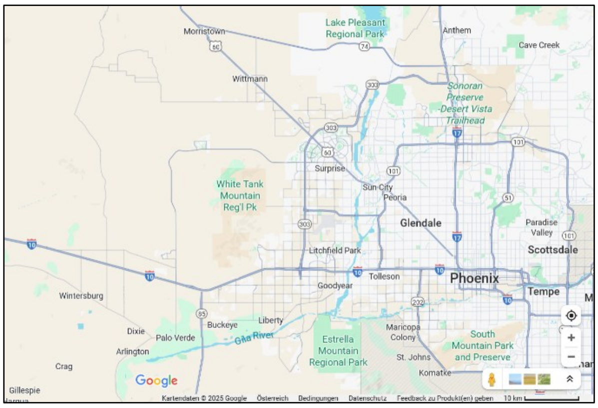

Figure 1. *Map of Tokyo on a 10km scale bar (Google Maps 2025).*
Figure 2. *Map of Phoenix, Arizona, on a 10km scale bar (Google Maps 2025).*

The map of Tokyo appears more chaotic, with a complex street network and extensive public transport overlays reflecting the city's high urban density and sheer size. In contrast, the visual structure of the map of Phoenix is more spacious and grid-like. It emphasizes major roads, suburban regions, and parks, which points towards a low-density, car-oriented urban structure. Cultural interpretations of colour also vary greatly; for instance, red may mean 'danger' in Western cultures, but represents luck and prosperity in Chinese culture. Additionally, languages may also require different placement strategies (right-to-left scripts vs. left-to-right scripts) (Zhunis et al. 2021: 2). Zejdlik & Vozenilek say that standardization efforts must therefore consider usability and regional representation, as inconsistencies in labels, analytical tools, and navigational features can significantly alter how users engage with and trust digital maps (2025: 2).

As the most widely used mapping platform, Google Maps has particularly shaped how people perceive spatial reality. Muehlenhaus warns that "if something does not exist on a Google Map, it will cease to exist in people's mental maps of places unfamiliar to them" (2014: 11). This gives way to possible omission bias and digital erasure, when places, roads, or even entire communities may become invisible simply because they are not represented on dominant mapping platforms. A striking recent instance is the renaming of the Gulf of Mexico (The White House 2025): following Executive Order 14172, signed by President Donald J. Trump on January 20, 2025, Google Maps began displaying 'Gulf of America' for US users, while showing 'Gulf of Mexico (Gulf of America)' for users elsewhere (see *Figure 3* and *Figure 4*). No other nation has officially adopted the new name, yet millions of users now see a different map depending on where they are located. This illustrates how the influence of dominant platforms like Google Maps on spatial perception goes beyond omission – it also extends to actively shaping which version of reality users see.

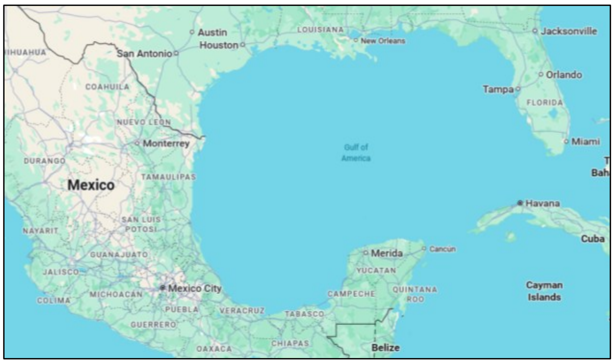

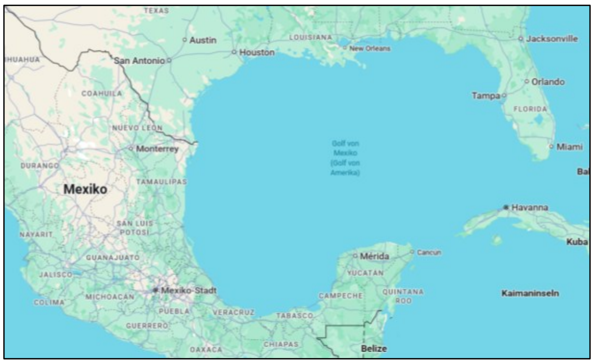

Figure 3. *The Gulf of America on Google Maps (US location) (2026).*
Figure 4. *The Gulf of Mexico on Google Maps (every other location) (2026).*

2.2. Commercialisation of Digital Maps

In the 2010s, open-source data and content was becoming widely recognized among GIS- and geo-data-users. Tools like QGIS provided free alternatives to commercial GIS software, which made spatial analysis more accessible. At the same time, mapping platforms transformed into powerful commercial ecosystems, driven by the explosion of big data. Platforms like Google Maps and Bing Maps no longer just display geographic information; they aggregate, analyse, and monetize spatial data (Google Maps Platforms 2024). This includes satellite imagery, street-level photography, and real-time user-generated data, all of which play a crucial role in the digital economy.

Companies leverage these vast data collections to optimize their businesses. According to Polaris Market Research (2024), “the digital map market was valued at approximately $25.28 billion [in 2023] and is projected to grow to $66.58 billion by 2032.” These numbers highlight how geospatial data has become a key driver of advertising revenue and targeted marketing, going far beyond its original function in navigation. A significant factor in this development is Google Maps' early introduction of its API (Application Programming Interface) in 2005, which solidified its dominance in the digital mapping industry. This decision allowed third-party developers to integrate Google’s maps into their own websites and applications, effectively transitioning Google Maps from a mapping service into a platform (Kenney & Zysman 2020: 7). By offering free access for affordable rates, Google Maps quickly became the world’s most used API by 2013, with over one million users (Google Platforms 2013). This strategy ensured that even Google’s competitors, like Yelp and Booking.com, have now come to rely on Google Maps’ mapping infrastructure. The introduction of billing requirements in 2018 transformed every embedded Google Maps API into a potential revenue source for Google, making even small businesses financially dependent on Google’s mapping services (Kenney & Zysman 2020:70). Additionally, this API integration enables continuous data collection, as Google receives user location data each time an embedded map is accessed. This means that Google Maps has two major revenue streams: Firstly, direct payments from businesses using Google Maps on their websites, and secondly, targeted advertising based on user interaction.

As a result, online maps have become interactive, two-way communication systems. Users actively and passively contribute data – for example through reviews, location tracking, traffic patterns, and search queries – which platforms, in turn, monetize via advertising, business rankings, and personalized recommendations. Muehlenhaus describes this shift, stating, “We are living in a data-rich society, which means we have to decide which data are, or better still, which information, is worth mapping. Digital data is everywhere and easy to create, get, find, steal and distribute” (2014:13). This unprecedented accessibility has blurred the line between cartographers and users, as our everyday interactions reshape digital maps. One main issue within this context is that with the proliferation of platforms such as ArcGIS, Mapbox, and HERE, anyone is allowed to create and share web maps, often without any formal cartographic expertise. Zejdlik & Vozenilek highlight that while this accessibility fosters innovation, it also raises concerns about data quality and adherence to cartographic principles (2025: 13). Their study notes that inconsistent symbology, misaligned navigation elements, and regional inaccuracies can create usability issues, particularly in web maps that serve as primary sources of geographic information for the public. These inconsistencies become especially problematic when users rely on these maps for navigation, urban planning, or decision-making. Even within official, accepted, and ‘trusted’ online maps, this interactivity comes not without risks.

Here, Critical GIS becomes particularly relevant, which developed in the late 1980s and 90s as a response to the claim that GIS is a purely neutral and technical representation of space (Sheppard 2005: 6). As already mentioned, early GIS was closely tied to governments and the military, which shaped both the availability of geospatial data and what it was used for. Critical GIS scholars challenged these assumptions by emphasising that GIS technologies are embedded in social, political, and economic contexts, and any decisions about how data is represented can reinforce unfair power dynamics and spatial inequalities (Sheppard 2005: 7, Thatcher et al. 2016: 821) These concerns remain highly relevant in the context of contemporary digital mapping platforms, where commercial actors increasingly control geospatial data. As users expect that maps are always up to date, it also means that platforms and websites have an incentive to constantly track our movements and analyse our behaviours. Concerns about privacy, surveillance, and data security are a hotly debated topic time and time again. Additionally, the filtering and ranking of geospatial information may lead to bias and selective visibility, where certain locations or businesses are displayed over others based on financial interests. In response to these critiques, scholars did not simply reject GIS but sought to reconstruct it. Elwood (2006: 693) traces how participatory GIS emerged from this impulse, seeking to embed GIS in more inclusive decision-making processes. However, she also cautioned that these reconstructions introduce their own contradictions, and that as participatory GIS becomes more mainstream, it does not always preserve its original critical commitments (2006: 700).

Still, maps and mapping practices can also function as tools for social transformation by shifting existing power relations. Pavlovskaya argues that Critical GIS can make alternative economic and social practices visible, particularly those that are marginalised or overlooked by dominant, profit-driven ways of understanding space (2018: 52). By visualising less visible forms of economic activity, new mapping can challenge established cartographic work and open up space for alternative interpretations of how places and economies function.

2.3. Artificial Intelligence in Digital Mapping

Beyond these issues, large-scale geospatial data is a prerequisite for training GeoAI models, with commercial platforms like Google Maps and Bing Maps emerging as particularly significant data sources (Kang, Gao & Roth 2024: 605). “Through data aggregation, big data supports the application of algorithmic analytics, popularly associated with machine learning and loosely labelled ‘AI’.” (Kenney & Zysman 2020:

58). One of the most promising research areas is called GeoAI (Gao 2021; Janowicz et al. 2020), which with the help of machine learning and deep learning, focuses on feature detection, real-time spatial analysis, and predictive modelling (Kang et al. 2024: 3).

However, while AI enhances map-making processes, it is also increasingly being explored for its potential in generating maps autonomously. By processing user inputs, such as text descriptions, images, or geodata, machine learning models can generate maps without requiring GIS-knowledge (Wu et al. 2024: 3). This development paves the way for autonomous GIS agents that can execute complex spatial analyses, retrieve relevant geodata, and are able to improve their outputs with minimal human intervention (Li et al. 2025: 505). A notable example is the LLM-Geo agent, which translates natural language queries into executable geoprocessing workflows and even debugs its own code when needed (Li et al. 2025: 513). Generative AI also supports customization, for example by enabling users to specify map styles, scales, etc.; as well as facilitating spatial simulation, which can help users visualize hypothetical scenarios better (Wu et al 2024: 2). In summary, the main aspect that distinguishes GeoAI from generative AI, is its reliance on geographic knowledge and domain-specific insights, rather than solely generating new map-related content.

However, there are also several risks that become particularly relevant with GeoAI systems. The GeoAI vision paper by Mai et al. emphasizes that foundation models for GeoAI may suffer from so-called ‘geographic hallucinations’ and ‘geographic bias’, which means that spatially inaccurate or misleading outputs – despite appearing plausible – are being produced (2024: 33). For example, a location might be misrepresented or systemically favoured. Such errors often stem from poor training data and population-biased ranking statistics. ‘Temporal bias’ is another aspect to factor in, which is the tendency of foundation models to rely on present-day knowledge and overlook historical context, which leads to errors when geographic names or places have changed over time (2024: 32). Lastly, this all ties together with serious ethical concerns, because AI-driven spatial systems translate geodata into automated decisions, which in turn can have direct economic, social, and physical consequences. As Chen et al. show in their newest study, many ethical issues associated with GeoAI are fundamentally data-driven and originate across the entire

GIS workflow (2025: 2). When GeoAI systems rely on outdated, incomplete, or biased geodata, they can produce misleading predictions, leading to economic harm in fields such as finance, healthcare, urban development, or security, where spatial assessments are used to allocate resources or assess risk (2025: 18). At the same time, because geodata contains sensitive location information – such as home addresses – by default, GeoAI applications can enable privacy violations and intrusive forms of surveillance. Ethical risks are further intensified by uneven data coverage, as certain regions or social groups may be systematically underrepresented, which reinforces already existing spatial inequalities and social injustice. From an environmental and ecological perspective, inappropriate use of geospatial data may also contribute to overexploitation or environmentally harmful decision-making (2025: 18). Taken together, these examples illustrate that ethical challenges in GeoAI are not limited to technical inaccuracies but are closely tied to how spatial data is produced and processed.

Furthermore, while AI is transforming cartographic work through automation and advanced spatial analysis, there is also a third major shift occurring: Artificial Intelligence is increasingly bypassing maps altogether, by instead providing users with ranked lists and recommendations. Whether searching for a nearby café, a hotel, or the best route to a destination, users often solely rely on AI-based rankings. Ravi & Vairavasundaram (2016) explored in a study how recommender systems in travel services have evolved, particularly with the integration of social network data in order to improve prediction accuracy. They concluded that as AI-driven filtering techniques are becoming more advanced, recommendation systems prioritize personalization over direct spatial interaction. While recommendation lists have been a common feature in digital services for some time, early systems relied on static data sets, user reviews, and sponsored content rather than dynamic AI-driven insights (Chen, Chen & Wang 2015: 100). More recent advancements in deep learning and real-time data processing now enable AI to rank and filter locations based on relevance, user preferences, and live information (Zhang, Lu & Jin 2021: 339) significantly reshaping local search, travel planning, and urban decision-making (Ravi 2016: 2). Results are consequently tailored to an individual's needs. For example, Google's Recommendations AI leverages deep learning models to provide highly personalized recommendations at scale, adapting dynamically to changing user behaviour and preferences (Google Cloud 2021). As a result, digital maps are becoming less of a primary interface and more of a background tool, supporting rather than guiding decision-making.

2.4. Dark patterns

The increasing reliance on algorithms in web maps certainly raises concerns about transparency and commercial influence. The presence of so-called "dark patterns" or “deceptive patterns” are manipulative user-interaction strategies that deceive users or prioritize business interests over user needs, are a rarely discussed phenomenon when dealing with maps. The term was first coined by Harry Brignull in 2010, who describes them as “a user interface that has been carefully crafted to trick users into doing things, such as buying insurance with their purchase or signing up for recurring bills” (DeceptivePatterns.org 2023). However, there is no universally agreed definition, even if in the past 15 years the term has garnered a lot of attention from businesses, consumers, government regulations, and researchers alike. One reason for the rise of deceptive patterns is the growing dominance of metrics-driven design: tracking user behaviour has become increasingly easy, while privacy concerns are often muted by the fact that users do not actively feel “seen” (Brignull 2023: 20). Combined with the ability of large-scale A/B testing, widespread data processing, and the easy replication of successful design patterns across platforms, manipulative interfaces have become more norm than exception (Brignull 2023: 22). A/B testing refers to the comparison of different interface variations to see which drives the most user engagement. On large platforms, these tests run at scale across millions of users in real time, which allows companies to quickly adopt new design choices if they boost performance (Young 2014: 2). As a result, patterns, regardless of what kind, can spread rapidly.

To combat these developments, many laws and regulations have sprung up the past years to help govern the use of deceptive patterns. The EU Digital Markets Act and Digital Services Act, which was implemented in 2022, and which aims to ensure greater accountability by regulating large online platforms, is arguably the furthest ahead when it comes to protecting the rights of users in the digital space. It defines dark patterns as:

> “[…] [online] practices that materially distort or impair, either on purpose or in effect, the ability of recipients of the service to make autonomous and informed choices or decisions. Those practices can be used to persuade the recipients

of the service to engage in unwanted behaviours or into undesired decisions which have negative consequences for them." (67)

There are significant law cases that are already in motion thanks to the EU Digital Markets and EU Digital Services Act. For example, the platform X failed to comply with the required transparency on dark patterns, which resulted in the European Commission now considering a potential fine of over one billion dollars (Tech Policy Press 2025). AliExpress is another famous example for violating the Digital Services Act and could face fines of up to 6% of its global annual turnover (FairPatterns 2024). In particular, the European Commission opened a formal investigation into AliExpress in March 2024, citing concerns over misleading interface designs that potentially exploit consumers, especially minors, and violate content moderation obligations under the Digital Services Act (Lomas 2024).

In e-commerce, various types of dark and deceptive patterns have been identified and classified, the most influential and foundational one being Harry Brignull's in 2010. However, as patterns continuously shift and evolve, there is no singular taxonomy, and recent research has come up with new categorizations in a variety of domains. For example, Gray et al. expanded Brignull's definitions to include more categories (2018), and Bösch et al. made a taxonomy for patterns invading privacy, specifically (2016). Lewis identified design patterns in mobile games (2014), and Greenberg et al. introduced deceptive patterns evolving around proxemic interactions (2014). That there are so many different classification systems for dark patterns demonstrates just how difficult it is to 'pin them down', as many patterns overlap functionally, only exist in certain fields of the online world, rely on subtle design cues, and always balance between ethics and business interest.

I will now only briefly introduce the latest classifications and definitions of Brignull, Leiser, Santos & Doshi from the deceptivepatterns.org website (2023), as they are some of the most widely cited researchers in this field, and all other existing typologies of deceptive patterns are connected to their work:

Comparison Prevention	The user struggles to compare products because features and prices are combined in a complex manner, or because essential information is hard to find.
Confirmshaming	The user is emotionally manipulated into doing something that they would not otherwise have done.
Disguised ads	The user mistakenly believes they are clicking on an interface element or native content, but it is actually a disguised advertisement.
Fake scarcity	The user is pressured into completing an action because they are presented with a fake indication of limited supply or popularity.
Fake Social Proof	The user is misled into believing a product is more popular or credible than it really is, because they were shown fake reviews, testimonials, or activity messages.
Fake Urgency	The user is pressured into completing an action because they are presented with a fake time limitation.
Forced Action	The user wants to do something, but they are required to do something undesirable in return.
Hard to Cancel	The user finds it easy to sign up or subscribe, but when they want to cancel, they find it very hard.
Hidden Costs	The user is enticed with a low advertised price. After investing time and effort, they discover unexpected fees and charges when they reach the checkout.
Hidden Subscription	The user is unknowingly enrolled in a recurring subscription or payment plan without clear disclosure or their explicit consent.
Nagging	The user tries to do something, but they are persistently interrupted by requests to do something else that may not be in their best interests.
Obstruction	The user is faced with barriers or hurdles, making it hard for them to complete their task or access information.
Preselection	The user is presented with a default option that has already been selected for them, in order to influence their decision-making.
Sneaking	The user is drawn into a transaction on false pretences, because pertinent information is hidden or delayed from being presented to them.
Trick Wording	The user is misled into taking an action, due to the presentation of confusing or misleading language.
Visual Interference	The user expects to see information presented in a clear and predictable way on the page, but it is hidden, obscured or disguised.

Table 1. *Types of Deceptive Patterns.* (Brignull et al. 2023).

While the types in *Table 1* are common examples for general e-commerce and subscription services, similar tactics can appear in web maps. Not all those patterns can be applied 1:1, but many of them are of relevance. Web maps represent a comparatively underexplored area, even though they combine factors such as UI design, algorithmic filtering, and commercial pressures that make them especially vulnerable to manipulative practices.

3. Methodology

To investigate how users perceive potentially deceptive patterns in online maps, a multi-step methodological approach was adopted. First, deceptive patterns in web maps were identified based on existing research and literature in e-commerce and marketing. Next, a survey was conducted to capture users' expectations and perceptions regarding the presence and role of such patterns in web maps. Importantly, the survey did not aim to determine whether these patterns were intentionally put there by platforms or to infer the technical motivations behind specific interface behaviours. Instead, it focused on how users experience and interpret map-based search interfaces, regardless of whether observed effects arise from deliberate design choices or from technical constraints such as spatial gridding or data aggregation. Based on the survey results, a cross-platform analysis was subsequently conducted across twelve different online portals employing map-based web search.

3.1. Deceptive Pattern Identification in Map-Based Web Search

The following list of deceptive patterns in map-based web search in this chapter is inspired by Brignull et al.'s established framework of deceptive patterns (*Table 1*) but adapts and combines certain categories to better reflect the specific dynamics of map-based search interfaces. For example, patterns such as *Confirmshaming*, *Nagging*, *Hard to Cancel*, and *Hidden Subscription* typically occur in subscription-based services or during transactions and were therefore excluded due to their limited relevance to interactive map interfaces. Likewise, categories such as *Trick Wording* and *Visual Interference* were reviewed but omitted, as they refer to specific linguistic or visual manipulations that are not prominently used in current map design. This thesis focuses only on observable and functionally relevant patterns within a spatial search environment. As of now, there is no published work on the types of deceptive patterns in maps, specifically, subsequently this thesis hopes to address this gap.

1. **Camouflaged Advertising – Disguised Ads/Sneaking**: While *Sneaking* and *Disguised Ads* are defined separately in Brignull's typology, both patterns revolve around the invisible or misleading integration of commercial content into what users assume to be neutral results. In the context of map-based search interfaces, these patterns often coincide or overlap: sponsored listings are inserted (sneaking) while looking indistinguishable from organic results (disguised ads). Users may click on pins thinking it's a genuine nearby option, not knowing it's a paid ad. Given their functional similarity and frequent co-occurrence in platforms like Apple Maps, Yelp, and Booking.com, this study combines them under the broader category of *Camouflaged Advertising*.

2. **Non-transparent Map Filtering**: This thesis also introduces a new category, which describes how relevant locations are selectively shown or hidden depending on zoom level in web maps, without any explanation of the underlying prioritization logic. While zoom-dependent generalization is a necessary aspect of cartographic design, the lack of transparency regarding *which* locations are prioritized and *why*, invites further questions. No bigger commercial online portal that uses map search offers any explanation to this process (as of May 2025). This category differs from Camouflaged Advertising in that the results may not be sponsored or paid; however, the result is similar: the user may believe that what is displayed is the only available and correct option, and additional listings remain hidden due to the map's zoom level or internal ranking logic.

3. **Fake Scarcity:** Some map platforms show urgent messages like *"Only 1 room left at this price"* or *"Booked 3 times today"*. In 2020, the accommodation search platforms Booking.com, Agoda, and Expedia were all accused of misleading consumers by the European Commission and national regulators for faking scarcity (Brussels Times 2024). Only as of November of 2024, changes have finally been made on the websites to fall in line with EU law (Booking Holdings Inc.'s Digital Markets Act Compliance Report 2024).

4. **Fake Social Proof:** Listings may include artificial reviews, or reviews that lack transparency, and automated *“X people are viewing this”* messages. There have already been public cases of proven fake reviews, for instance on TripAdvisor. In 2017, a journalist set up a fake restaurant in London using fake reviews as an experiment, and in a span of 6 months was listed as one of the top restaurants on TripAdvisor (Vice 2017). This is problematic insofar as the number of reviews is shown when hovering with the mouse over a pin in the TripAdvisor map. A study done by Nielsen in 2015, which, while predating more recent concerns around AI-generated and fake reviews, reported that over 60% of people trusted online reviews most after personal recommendations (2015: 13).

5. **Fake Urgency**: Time-limited offers shown on location pin cards, or phrases like “deal ending soon” push users to act more quickly. Similar to Fake Scarcity, multiple accommodation portals like Agoda were accused of doing this in 2020.

6. **Hidden Costs:** Map listings don’t show all fees upfront (e.g. taxes or resort fees are only revealed at the check-out). A prominent case study is Airbnb, which did – and still does – not showcase the full price on their maps, depending on the country. At the end of the booking process, an additional cleaning and service fee is added, which can in some cases, add to over a third of the original price. Several national regulators have stepped in to stop this deceptive pattern, such as the Australian Competition and Consumer Commission (ACCC) in Australia, or the Norwegian Consumer Authority and European Commission in Norway (Brignull 2023: 96). In those countries, Airbnb complied and has changed its ways, however it has not stopped doing so in many other countries (Japan, USA, etc.).

7. **Obstruction:** Many map interfaces reduce on-screen information out of necessity to make maps visually cleaner and less cluttered. However, in some cases, this reduction may also serve to intentionally slow down user interaction. When essential filters or features are buried behind multiple steps, or when users must click on each individual pin to access full details – rather than being offered a unified overview – this becomes problematic. For

example, on Kayak.com's "Explore" map version of searching for flights, selecting a single flight is not possible, only a roundtrip. Only after selecting the dates and flight and being forwarded to the list version of offers, the setting is changeable. When searching for properties in Vienna on Sotheby's International Realty, the platform presents a map alongside a list of available listings. However, the map displays only a small number of pins, like two or three, even though a bubble on the map itself states that dozens of properties (e.g., "47 available") match the search.

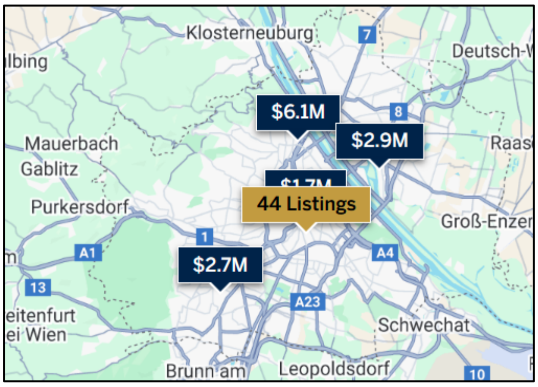

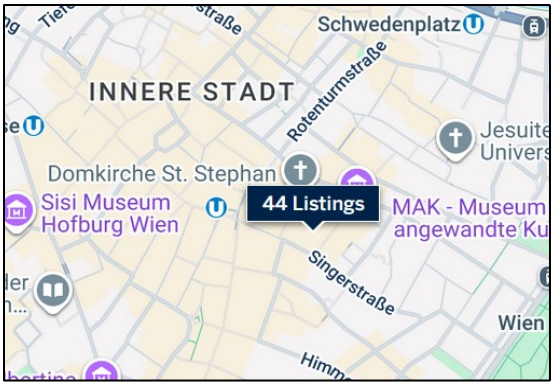

Figure 5. *Property listings on a zoomed-out map of Vienna on Sotheby's Realtor International.*
Figure 6. *Property listings on a zoomed-in map of Vienna on Sotheby's Realtor International.*

It should be noted that this may partly reflect individual property owners' preferences to withhold exact locations. Nonetheless, as seen in *Figures 5* and *6*, the map does not update when zooming in, nor does it reveal additional listings. To access property details or see further options, users must engage with the list view beside the map. Even after clicking a specific listing, the platform does not highlight its location on the map, nor allow for spatial comparison with others. Instead, users are redirected to an individual listing page, where the property's location is buried at the bottom behind an "expand map" feature. At this stage, the map displays only the selected property, but prevents any spatial comparison with alternative listings. As such, the whole map interface is largely non-functional for location-based search, offering no overview or context.

8. **Preselection:** This pattern refers to when filters or sorting (e.g., "recommended" or "sponsored first") are auto-applied. This is common across most commercial web map interfaces. Usually, after entering a search query, a default setting is applied called 'recommended' or 'best match', with either

hidden or no transparency on what ‘best’ means. For example, Booking.com offers some explanation to their ranking system via a “*Learn more*” link, which opens a separate tab. There, it says that it is a combination of sponsored placement, user reviews, ratings, click-through rate, gross, and net bookings, etc. However, this option is only available on the (list view) home page, not when searching with the map. Google Maps includes a tiny information icon with a short statement about how hotels and accommodations may be ranked based on the users’ previous activity, but not how anything else is ranked. Apple Maps, in contrast, offers no information at all. A further example of Preselection occurs on Booking.com, where the map-based accommodation search initially shows both available and unavailable properties. Only when manually clicking on a button on the side bar “show only available properties”, this will change. This raises the question of why non-available accommodations show up as map pins in the first place, when intuitively, the map-based search is only for available accommodations (Janowicz et al. 2024: 3). Lastly, a note has to be made about Non-transparent Map Filtering: the effect is similar to Preselection, as both display a default for users, but it differs in that it operates without any controllable user setting. Search results are based on unknown criteria, which makes it an algorithmic question rather than standard bias.

Another potential but unconfirmed pattern is *Spatial Anchoring*, where the map’s default viewport consistently emphasizes higher-priced or sponsored zones. While this behaviour may stem from technical or usability considerations, it raises questions about how initial spatial framing might influence user decisions and whether it serves commercial interests. However, no clear evidence of manipulative intent was found in the current research, hence it was excluded from this work, but the pattern may warrant future investigation.

To summarize, *Table 2* below lists all identified deceptive patterns in map-based web interfaces again:

Camouflaged Advertising	Sponsored content is inserted into the map in ways that visually resemble organic location results.
Non-transparent Map Filtering	Certain listings are only revealed at closer zoom levels, without explanation.
Fake Scarcity	When clicking or hovering over a map pin, messages like "Only 1 left" are shown to pressure users into fast decisions.
Fake Social Proof	Map pins or previews show vague popularity cues or false ratings.
Fake Urgency	Hovering or clicking on a map location triggers countdowns or suggest artificial time pressure.
Hidden Costs	The total price is only revealed after several clicks or navigating through external pages, while initial map results display misleading base prices.
Obstruction	Key map filters or settings are buried behind multiple clicks, hindering users from refining or adjusting the map view easily.
Preselection	Default sorting or filters are applied to map results without user input.

Table 2. *Types of Deceptive Patterns in Commercial Web Maps*.

Together, these eight patterns reveal that while many deceptive strategies used in map-based search stem from broader digital design practices, they are also shaped by spatial interaction and cartographic design choices. Taken this into consideration, some dark patterns may arise unintentionally due to technological and design limitations.

Peterson explains that early web maps relied on raster file formats, which became obsolete with higher-resolution screens (2014: 265). Vector-based formats improve scalability and adaptability, but require more processing power, sometimes causing delays or incomplete rendering. Similarly, Zejdlik & Vozenilek note that raster tiles (e.g., used by OpenStreetMap and Mapy.cz) limit label flexibility, leading to overlaps or unclear text, while vector tiles (e.g., used by Google Maps and Bing Maps) lead to more dynamic adjustments (2025: 13). Importantly for this thesis, vector tiles also enable personalization of map content, meaning different users may see different results, which opens the door for targeted advertising. Additionally, vector tiles improve readability but also come with inconsistencies at different zoom levels. These limitations may unintentionally obscure spatial information. So while these constraints are technical in nature, they may inadvertently produce effects similar to deceptive patterns — particularly in the case of Non-transparent Map Filtering, where visibility is uneven and unpredictable across zoom levels. Generally, it is important to note that “deceptive patterns are not always the result of rigorous research and careful craftsmanship. Sometimes they are just profitable accidents” (Brignull 2023: 14). According to Greenberg et al., such occurrences could also be classified as anti-patterns, which indicate a design failure or non-solution, but are not harmful on purpose (2014: 524).

3.2. Survey

To investigate people’s perceptions of deceptive design in digital maps, a five-point Likert-scale survey was developed, with the answer scale ranging from ‘strongly disagree’ to ‘strongly agree.’ In addition, demographic questions (age, gender) were included to allow basic subgroup analysis. The questionnaire comprised ten statements, eight of which were aligned with the previously identified deceptive patterns in map-based web interfaces. However, since deceptive patterns often operate subtly or only in interaction with other design elements (e.g., hidden costs only become apparent at checkout), it proved challenging to align all survey statements without making the survey too long or too complicated. Capturing such dynamics in the format of a single Likert-scale statement required a certain degree of simplification, where some contextual nuance might be lost. Furthermore, some patterns overlap conceptually (e.g., urgency and scarcity), making it difficult to isolate them without influencing how respondents interpreted the questions. To address this, the survey design aimed for statements that reflected the experience of users rather than the exact technical implementation of the patterns.

In order to ensure contextual accuracy and increase the validity of responses, three slightly different versions of the questionnaire were created, each tailored to a specific domain: (1) Hotels & Accommodations, (2) Points of Interest & Recommendations, and (3) Real Estate. While the majority of statements were identical across versions, minor wording adjustments were made where necessary to match the semantic and functional context of each category. For example, the statement *"Map markers should contain immediately visible key information like pricing"* was included in the questionnaires for Hotels & Accommodations and Real Estate, where price is a critical factor. However, for POIs & Recommendations, the statement was adapted to *"Map markers should contain immediate visible key information like type of venue and popularity"*, since pricing is typically irrelevant in that context (e.g., for public parks or viewpoints). All three versions of the questionnaire were then distributed digitally via a Google Forms link. Participation was voluntary, and responses were collected anonymously.

3.3. Cross-Analysis

The second part of this thesis involved an examination of multiple online portals that implement map-based web search as a main feature of their websites. Since platforms use online maps for different purposes, and ergo, the purpose of potentially deceptive patterns may differ, three broader categories were created to be able to give more accurate insights. The categories and examined websites are:

1. **Hotels and Accommodations:** Booking.com, Airbnb, Agoda, Expedia.
2. **POI Recommendations and Reviews**: Yelp.com, TripAdvisor, Bing Maps, Google Reviews.
3. **Real Estate**: Immoscout24, Zillow, Sotheby's International Realty, Rightmove.

Some of the previously identified patterns, especially Fake Social Proof, Urgency, and Scarcity, are inherently difficult to verify, since there is no way for the standard researcher to access the authenticity of reviews, the actual number of bookings, and other availability figures. For this reason, this study does not attempt to fact-check the truthfulness of such claims and rather focuses on the presence and prominence of those patterns. It assumes that user perceptions – and thus potential user manipulation – are shaped by appearance rather than fact. In practice, this means that the presence of cues such as star ratings or urgency messages was coded, without evaluating their factual accuracy. This approach aligns with EU regulatory reasoning, which says that "A commercial practice shall be regarded as misleading if it contains false information and is therefore untruthful or in any way, including overall presentation, deceives or is likely to deceive the average consumer, even if the information is factually correct […]" (Directive 2005/29/EC, Art. 6[1]).

Building on this definition, the coding process examined how potentially deceptive patterns were implemented and visible to users across the selected platforms. Each pattern was coded according to its clarity, labelling, and effect on user perception in both map and list views. A value of 1 (present) was given when a design element could reasonably mislead users during normal browsing, and 0 (absent) when transparency was achieved through clear labelling or accessible disclosure. In addition to this binary coding, qualitative notes were taken in special cases, when prominent aspects were notable. These notes provide further examples and are discussed later in the results section. For platforms offering both map and list views, the two were coded separately to capture possible differences in filtering, ranking, and user control.

Ranking systems labelled as "best match", "recommended" or similar were treated as a form of Preselection, because they reorder results in ways that are not transparent to users. Whether these rankings are AI-driven or based on conventional algorithms is usually not disclosed, yet both produce the same effect – appearing neutral while prioritizing certain listings for commercial or behavioural reasons.

Camouflaged Advertising was coded as present (1) when sponsored or promoted listings were not clearly labelled as such within the visible listing (in either map or list view). Platforms that disclosed paid placements only in hidden help pages, terms and conditions, or info buttons were still coded as camouflaged, since typical users would not encounter this information during normal use. Only when a clear, in-situ label such as *"Ad"* or *"Sponsored"* appeared directly on the listing was the pattern coded as absent (0).

Fake Urgency (availability opacity) was coded as present (1) when platforms displayed misleading signals of limited availability, though the form of this pattern differed across categories. For hotel and accommodation platforms, it referred to explicit urgency messages such as "only 1 room left," "booked five times today," or similar prompts that create artificial time pressure. For real estate platforms, the same logic was applied to the display of reserved or unavailable listings that remain visible among active results, suggesting competition or scarcity even when no transaction is possible. In contrast, this pattern was excluded for POI and review sites, since urgency mechanisms are typically completely absent there. Platforms that clearly labelled or separated unavailable listings, or avoided misleading urgency prompts altogether, were coded as absent (0).

As is common in exploratory interface audits, the coding was conducted by a single researcher using pre-defined criteria applied consistently across all platforms. While this approach is appropriate for a first systematic examination of this kind, some coding decisions, particularly for patterns like Fake Social Proof and Non-transparent Map Filtering, inevitably involved interpretive judgment. The results should therefore be read as indicative rather than definitive.

4. Results

4.1. Survey Results

The survey yielded up to 50 responses in each category, with several participants choosing to complete more than one set of items. This resulted in 50 responses for points of interest (POIs) and recommendations, 46 responses for hotels and accommodations, and 45 responses for real estate. Respondents varied in age from teenagers to individuals over 60 years old; however, in all three categories, the majority were young adults between 18 and 29 years of age, who accounted for approximately 70% (n ≈ 32–35 per category) of the total responses.

The questionnaires consisted of eight items addressing deceptive patterns in map-based interfaces, and two more general items relating to the map view vs. list view discourse ("*The map interface is relevant to me even when I already know where I want to go"* and *"I prefer browsing for points of interest using a map rather than with a ranked list"*). In addition, category-specific questions were included for hotels *("I trust urgency messages [e.g., 'Only 1 room left'] to appear only when there is a real time limit")* and for real estate *("I expect listings marked as 'reserved' to be shown separately from available properties"*). The following table presents the results of these items, reported as mean values and standard deviations on a five-point Likert scale, alongside the proportions of respondents who were in agreement (scores of 4–5) or disagreement (scores of 1–2). The exact questionnaire item wordings were abbreviated. Because responses were recorded on a 5-point Likert scale, the average responses were interpreted together with how widely opinions differed and how the answers were distributed. An increased standard deviation was read as greater heterogeneity, hence lower certainty, about the presence of a pattern. Consequently, inferences were based primarily on the mode, median, and the share of respondents selecting ≥ 4. The mean was treated as a secondary, comparative descriptor for cross-item comparison. ‘Strong perceived use’ was classified where the mode is ≥ 4 with ≥ 60 % agreement and either SD ≤ 0.80 or the median ≥ 4. ‘Moderate perceived use’ was applied to 50–59 % agreement or mode = 4 with SD up to ≈ 1.10. Values outside these ranges were labelled weak/indeterminate perception. When responses showed high variation (SD > 1.10) or two clear peaks, it suggests that participants were divided in their opinions. These cases were interpreted with caution, as they reduce the reliability of any claim.

Survey statements (shortened)	**Deceptive Pattern**	**Mean (M)**	**Standard Deviation (SD)**	**% of Agreement**	**% of Disagreement**
Maps should show all results (available & unavailable)	Fake Scarcity	3,26	1,28	54%	40%
Click/hover map pins to see immediate reviews & details	Obstruction	4,32	0,66	91%	1%
Ratings & reviews should be transparent	Fake Social Proof	4,32	0,77	89%	4%
Urgency messages can be trusted	Fake Urgency – category Hotels & accommodation	3,26	1,39	52%	39%
Reserved listings should be shown separately	Fake Urgency – category Real Estate	3,91	1,02	78%	13%
Map markers should contain visible key information like pricing.	Hidden Costs	4,18	0,81	84%	4%
Map should display all results regardless of zoom	Non-transparent Filtering	3,51	1,11	56%	23%
Sponsored results should be clearly marked	Camouflaged Advertising	4,12	1,01	75%	9%
Usually accept default recommendations	Preselection	3,09	1,13	46%	38%
Prefer browsing with map vs. list	---	3,78	1,04	67%	17%
Map is relevant even if I already know where to stay	---	3,95	0,94	79%	11%

Table 3. *Summary of Survey Results.*

As displayed in *Table 3*, responses varied considerably across the items, each of which was tied to a deceptive pattern, though some addressed these patterns more directly than others.

The most consistently positively received items all related to transparency and visibility of information, with agreement rates ranging from 75% to 91%. "*Click/hover map pins to see reviews & details*" (Obstruction) and *Ratings & reviews should be transparent* (Fake Social Proof) both produced mean scores of 4.32 (SD = 0.65 and 0.77), with agreement rates above 89%. "*Map markers should contain visible key information"* (Hidden Costs) reached M = 4.18 (SD = 0.81; 84% agreement). Likewise, "*Sponsored results should be clearly marked* "(Camouflaged Advertising) scored M = 4.12 (SD = 1.01; 75% agreement). These findings highlight a strong demand for clarity and transparency in map-based interfaces.

Items relating to scarcity, urgency, filtering, and preselection produced more mixed responses. "*Maps should show all results" (available & unavailable)* (Fake Scarcity) received M = 3.26 (SD = 1.28), with 54% agreement and 40% disagreement. Here, certain differences can be seen based on the respective category, which will be discussed in category-specific subchapter. The hotel-specific urgency item "*Trust urgency messages only if real"* also had a mean of 3.26 (SD = 1.39), with 52% agreement and 39% disagreement, while the real estate–specific urgency item "*Reserved listings should be shown separately"* achieved somewhat higher endorsement (M = 3.91, SD = 1.02; 78% agreement), though disagreement was still present among 13% of respondents. The statement "*Map should display all results regardless of zoom"* (Non-transparent Filtering) reached M = 3.51 (SD = 1.11), with 56% agreement and 23% disagreement. Finally, "*Usually accept default recommendations"* (Preselection) showed the most divergence (M = 3.09, SD = 1.12), with 46% agreement and 38% disagreement.

The two general items showed moderate to strong support. "*Prefer browsing with map vs. list"* scored M = 3.78 (SD = 1.04), with 67% agreement and 17% disagreement, suggesting that maps are often considered the better tool for exploring results. "*Map is relevant even if I already know where to stay"* reached M = 3.95 (SD = 0.94), with 79% agreement. This indicates that maps are still seen as relevant beyond the search process itself, and are used for orientation even when the destination of choice is already known.

4.1.1. Category-specific results: Points of Interest (Survey)

In the Points of Interest (POI) and Recommendations category, participants expressed clear expectations for transparency and trustworthy information, with agreement rates above 85% for transparency-related items. The strongest agreement appeared for "*I expect ratings and reviews that are shown in digital maps to be transparent, without limiting how users can rate or what they can write*" (M = 4.26, SD = 0.76) and *"I expect to be able to click on/hover over location map pins to see reviews, photos, and more details*" (M = 4.26, SD = 0.66). Similarly, "*The map should display all available results regardless of zoom level, even if they are clustered or aggregated at smaller zoom levels"* (M = 4.28, SD = 1.13) and "*I expect sponsored results to be clearly marked and visually distinct from unsponsored ones on a map"* (M = 4.14, SD = 0.90) received high, though slightly more varied, agreement.

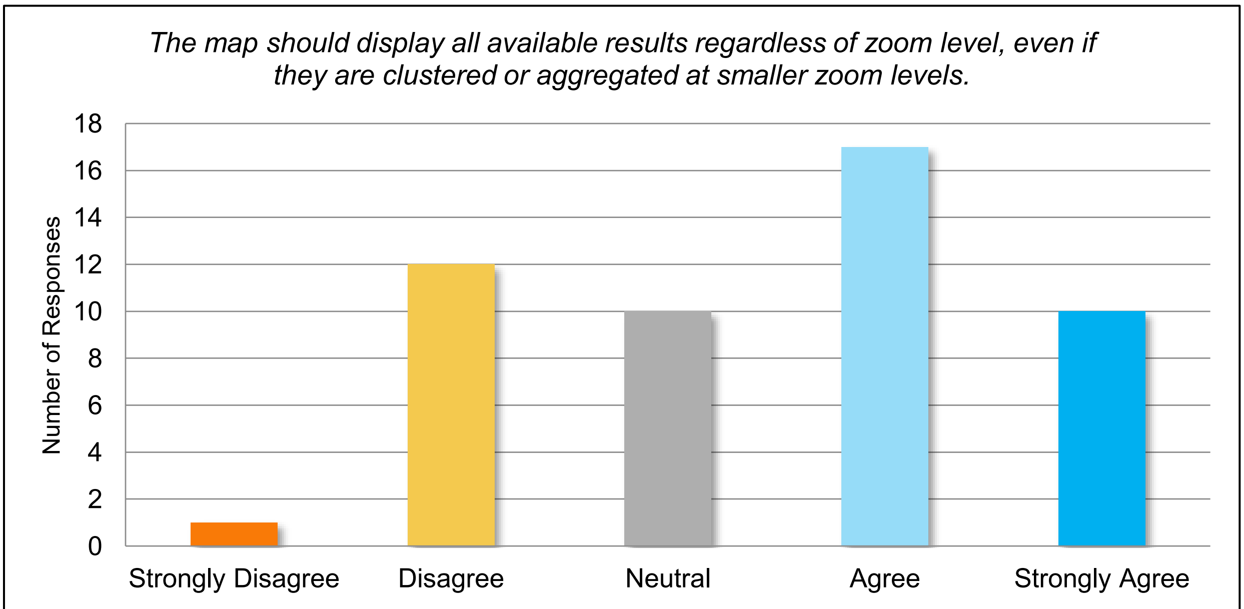

Figure 7. *POI user responses to "The map should display all available results regardless of zoom level."*

Figure 7 visualizes this distribution: most respondents agreed that maps should remain comprehensive across zoom levels, but the higher dispersion (visible in the bars for "Disagree" and "Strongly agree") points to differing opinions on how much aggregation is acceptable.

Moderate responses were given for "*Map markers should contain immediate visible key information like type of venue and popularity*" (M = 3.82, SD = 0.90), "*I prefer* *browsing for points of interest using a map rather than with a ranked list"* (M = 3.92, SD = 0.92), and "*The map interface is relevant to me even when I already know where I want to go"* (M = 3.46, SD = 0.99). These results suggest that map views are valued for local exploration, though not everyone relies on them as their main decision tool.

Greater variation appeared for "*When searching for points of interest, a map should display all results, i.e. open as well as closed places*" (M = 3.64, SD = 1.06), indicating differing expectations about completeness. "*I usually accept the default recommendations in maps without changing the sorting*" (M = 3.82, SD = 1.13) also showed divided opinions about automated Preselection. This category stands out because POI searches often involve spontaneous, everyday decisions – users prioritize quick access to reliable reviews over completeness or heavy filtering. As shown in *Figure 8*, responses clustered around "Agree" and "Disagree," indicating a clear split between users who rely on algorithmic sorting and those who consciously select filters.

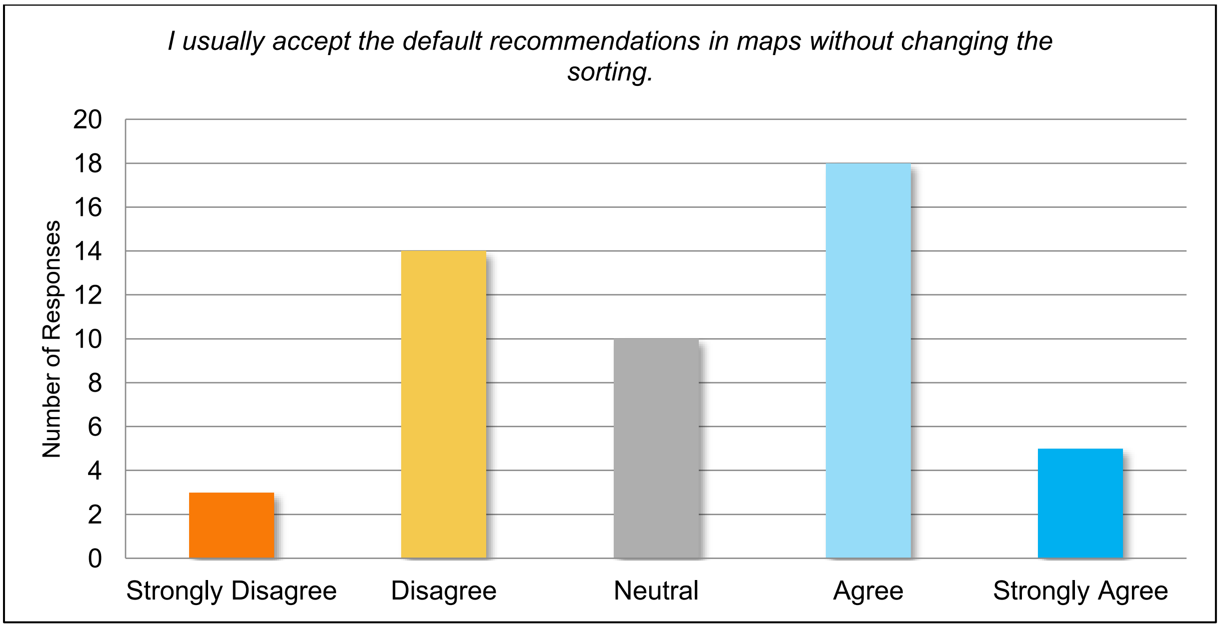

Figure 8. *POI user responses to "I usually accept the default recommendations in maps without changing the sorting."*

Overall, users agree strongly on the importance of transparency and detail, but differ in how much automation and completeness they expect from POI maps.

4.1.2. Category-specific results: Hotels & Accommodations (Survey)

For Hotels & Accommodation, survey takers again emphasized the importance of clarity and transparency. The highest mean values were found for *“I expect to be able to click on/hover over location map pins to see reviews, photos, and more details*” (M = 4.41, SD = 0.62) and “*Map markers should contain immediate visible key information like pricing*” (M = 4.33, SD = 0.85). Respondents also strongly agreed with *“I expect ratings and reviews that are shown in digital maps to be transparent, without limiting how users can rate or what they can write”* (M = 4.26, SD = 0.62). Labelling expectations remained high for “*I expect sponsored results to be clearly marked and visually distinct from unsponsored ones on a map*” (M = 4.04, SD = 0.79).

However, only moderate agreement appeared for “*The map interface is relevant to me even when I already know where I want to stay*” (M = 3.93, SD = 1.15) and “*I prefer browsing for accommodations using a map rather than with a ranked list*” (M = 3.85, SD = 1.15), showing that map views are appreciated but not essential for every user. Opinions became more divided for “*The map should display all available results regardless of zoom level, even if they are clustered or aggregated at smaller zoom levels”* (M = 3.59, SD = 0.85) and “*When searching for accommodations, a map should display all results, i.e. available as well as unavailable ones*” (M = 3.17, SD = 1.42). The item “I *trust urgency messages (e.g., ‘Only 1 room left’) to appear only when there is a real time limit”* (M = 3.26, SD = 1.05) also showed clear division, while *“I usually accept the default recommendations in maps without changing the sorting*” (M = 3.07, SD = 0.79) had the lowest overall agreement.

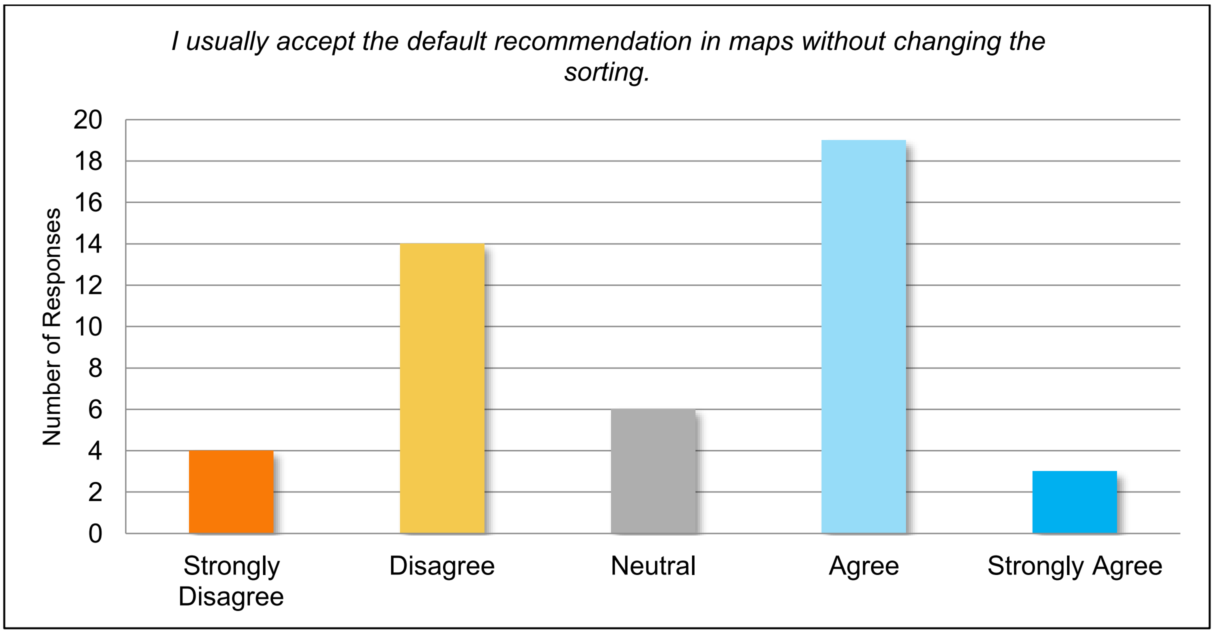

Figure 9. *Hotels & Accommodations user responses to "I usually accept the default recommendations in maps without changing the sorting."*

Figure 9 highlights a clear divide between users who rely on default rankings and those who deliberately change them, even if there is a slight leaning towards an acceptance of first results. This reveals very contrasting attitudes toward algorithmic guidance in booking interfaces.

This category is also distinct in that it has more commercial elements than Real Estate or Points of Interest, such as urgency messages, availability filters, and ranking algorithms, that directly affect trust. The mixed responses indicate that while users appreciate transparency, many remain cautious of persuasive booking cues and automated Preselection. Overall, participants value open and informative map displays, but their opinions differ widely in how much they trust system-generated filters and urgency indicators.

4.1.3. Category-specific results: Real Estate (Survey)

Once again, like in the other two categories, participants in the Real Estate section also showed straightforward expectations for clarity and transparency in map-based search. The highest mean values were found for "*Map markers should contain immediate visible key information like pricing*" (M = 4.42, SD = 0.69) and "*The map should display all available results regardless of zoom level, even if they are clustered or aggregated at smaller zoom levels*" (M = 4.42, SD = 0.69). Respondents also agreed strongly with *"I expect ratings and reviews that are shown in digital maps* *to be transparent, without limiting how users can rate or what they can write*” (M = 4.29, SD = 0.69) and “*I expect to be able to click on/hover over location map pins to see reviews, photos, and more details*” (M = 4.29, SD = 0.69). Expectations for labelling were slightly higher in this category than in others, as shown by “*I expect sponsored results to be clearly marked and visually distinct from unsponsored ones on a map”* (M = 4.44, SD = 0.55).

Opinions were more mixed regarding overall usability and filtering. The statement *"I usually accept the default recommendations in maps without changing the sorting"* (M = 3.04, SD = 1.11) showed the most divided responses of all items, with 44% agreement and 42% disagreement, which closely mirrors results from the other two categories. For instance, “*I prefer browsing for real estate using a map rather than with a ranked list*” (M = 3.49, SD = 1.06) and *“The map interface is relevant to me when I already know the area where I want to live in*” (M = 3.49, SD = 1.06) show that while maps are appreciated, not everyone relies on them equally. Finally, “*I expect listings marked as ‘reserved’ to be shown separately from available properties*” (M = 4.18, SD = 1.09) and “*When searching for real estate, a map should display all results, i.e. available as well as unavailable results*” (M = 2.91, SD = 1.28) reveal considerable differences in opinion about how much completeness users prefer in real estate maps.

Overall, the results show broad agreement on transparency-related features such as pricing visibility and labelling, with agreement rates consistently above 75%, while preferences diverge more on completeness and how much filtering or automation users want when browsing real estate, with agreement levels ranging between 45% and 60%.

4.2. Analysis Results

This chapter presents a comparative analysis of deceptive design patterns across the examined map-based platforms. The goal was to identify how these patterns differ between platform categories and how they relate to user expectations observed in the survey.

Based on the survey results, two patterns were excluded from the analysis: Obstruction and Hidden Costs. While Obstruction is a recognized category in dark pattern literature, it proved too broad and context-dependent to be analysed consistently across map-based search interfaces. As mentioned in previous chapters, obstruction typically refers to additional ‘friction’ in user experience, i.e. excessive clicks or hidden settings. However, it varies widely depending on the task and device. In the survey, the related statement *“I expect to be able to click on/hover over locations to see reviews, photos, and more details”* was predominantly answered with strong agreement, which reflects a universal expectation and means that users see it as a standard function rather than a deceptive technique. Similarly, Hidden Costs was excluded, because this pattern usually appears beyond the search interface at the checkout and varies by country regulation, taxes, user account, timing etc. To verify and check, it would require completing transactions, but since this thesis is limited to a map-based search experience, any results would be unreliable and not comparable across platforms.

The following section reports (i) overall prevalence of coded patterns across platforms, (ii) differences between map and list views, and (iii) category-specific mechanisms (Hotels & Accommodations; POI & Reviews; Real Estate).

	Hotels & Accommodations				POIs				Real Estate			
	Agoda	AirBnB	Booking.com	Expedia	Bing Maps	Google Reviews	TripAdvisor	Yelp	immoscout24.at	Rightmove	Sotheby's International Realty	Zillow
Camouflaged Advertising	1	0	0	1	1	1	1	1	1	1	1	1
Fake Scarcity	1	0	1	0	0	0	0	0	0	0	0	0
Fake Social Proof	1	1	1	1	1	1	1	1	0	1	0	0
Fake Urgency	0	0	1	0	0	0	0	0	1	0	1	0
Non-transparent Map Filtering	1	1	1	1	1	1	1	1	1	0	0	1
Preselection	1	1	1	1	1	1	1	1	1	0	1	1

Table 4. *Presence matrix of deceptive patterns in platform maps.*

The presence matrix across all platforms shows a clear baseline: Preselection and Non-transparent Map Filtering appear on almost every site, suggesting that visibility is shaped by default ranking and is constrained to the current map view. Camouflaged Advertising is common (though labelling clarity varies by platform), and Fake Social Proof is widespread wherever reviews are integrated. By contrast, Fake Scarcity and Fake Urgency are category-specific: urgency/scarcity cues surface only in Hotels & Accommodations (e.g., “only X left”), most notably in Agoda and Booking.com – see *Figure 10* – while Real Estate expresses urgency through the continued display of reserved/sold items among active results. POI platforms do not use these pressure cues.

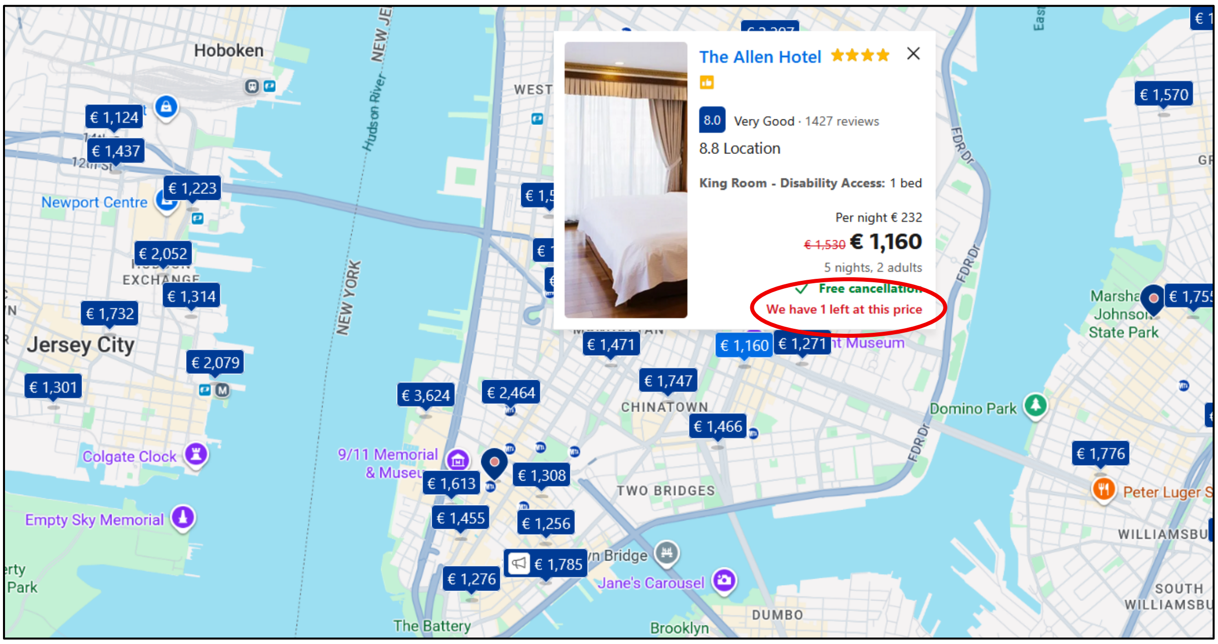

Figure 10. *Urgency cue ("We have 1 left at this price") displayed on the Booking.com map interface when searching for a hotel in New York.*

Across platforms, a few outliers stand out. Rightmove is the only service in the set that does not show evidence of Preselection or results tied to the current map view (Non-transparent Map Filtering), making it the main exception to the baseline. This is because it has no 'recommended', or 'best match' filter, it only shows newest to oldest by default. All other platforms showed evidence of this filtering effect, where additional listings only became visible after zooming in or panning the map, without any notice that the initial results were incomplete. An exception is Sotheby's International Realty, which displays bubble counts and reveals the full set of properties as users zoom, and was therefore coded 0 for this pattern.

While other platforms like Airbnb have previously used bubble counts, their interface appears to have changed over the course of writing this thesis – likely due to ongoing A/B testing – and no longer displays them consistently. As a result, it remains unclear which listings are shown at any given zoom level and why, with many options staying hidden without explanation. *Figures 11* and *12* demonstrate this – at first look, there is only one listing in the centre of Cape town, however, a zoom reveals additional listings, with more unnumbered bubbles appearing.

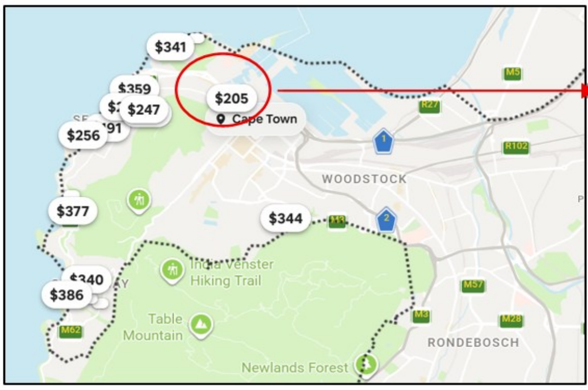

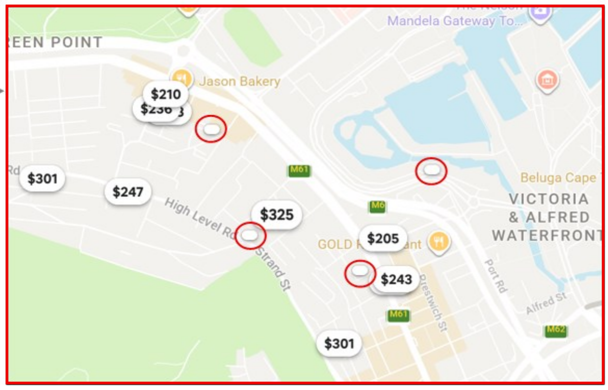

Figure 11. *Airbnb map interface of Cape Town zoomed out (left).*
Figure 12. *Airbnb map interface of Cape Town zoomed in (right), with bubbles appearing that indicate more available spots.*

Within Hotels & Accommodations, Airbnb and Booking.com had clear in-situ labels for promoted results (Camouflaged Advertising), whereas Agoda and Expedia did not; sponsored offers were visually indistinguishable from normal ones and were only revealed as ads when hovering over the listing or inspecting the link source. Pressure cues inside the hotel category are uneven: Fake Scarcity appears on some providers (e.g., Agoda, Booking.com) but not others (Airbnb, Expedia), and Fake Urgency is limited to Booking.com where countdown timers (“deal ends in X hours”) sometimes appear, but naturally, there is no way to verify if it really is a high demand period. These contrasts reveal that the use of deceptive patterns is not uniform but adapted to each platform’s interface design and commercial model.

When looking at *Table 4*, the distribution of deceptive patterns highlights that deceptive mechanisms in map-based interfaces primarily manipulate visibility and ranking logic, rather than using overt persuasive pressure, which are commonly found in normal e-commerce interfaces. In most cases, the user’s perception of “available options” is conditioned by pre-applied filters or by result clustering within the visible map area. Such mechanisms effectively narrow choice but maintain the *appearance* of an open search space. In practice, this means that when using map-based search, one is often interacting with a curated subset of listings that reinforce commercial priorities of the platform, for example by giving more visibility to sponsored or popular listings. In contrast, the selective use of urgency and scarcity strategies shows that very visible deceptive patterns are not a defining feature of map-based interfaces. These pressure tactics rely heavily on time-sensitive contexts (e.g., booking, limited inventory) and therefore align more with transaction-oriented domains. Their absence from POI and Real Estate maps indicates a prioritization of trust and informational reliability over short-term persuasion, at least in these environments.

Apart from preselection, Camouflaged Advertising was the second deceptive pattern that had a very strong presence across nearly all platforms. While the overall prevalence is high, the degree of visual disclosure varies significantly depending on how the advertising is embedded in each interface. Tripadvisor promotes listings through “Sponsored Placements,” which appear above natural results and within destination pages, while Expedia allows hotels to bid for visibility through its “TravelAds Sponsored Listings,” positioning paid properties in standard search results. It actually provides one of the clearest public explanations of how sorting and filtering affect ranking visibility, yet this transparency is unfortunately limited to the documentation level: in the map view, ads remain indistinguishable from unsponsored results. Yelp similarly prioritizes advertisers’ businesses, openly acknowledging in its support page that sponsored results appear before organic ones. However, as can be seen in *Figure 13*, they are not visually differentiated in the map interface, where all locations are represented by numbered pins. Hovering over the number – see *Figure 14* – does also not reveal further clues except the location name and rating.

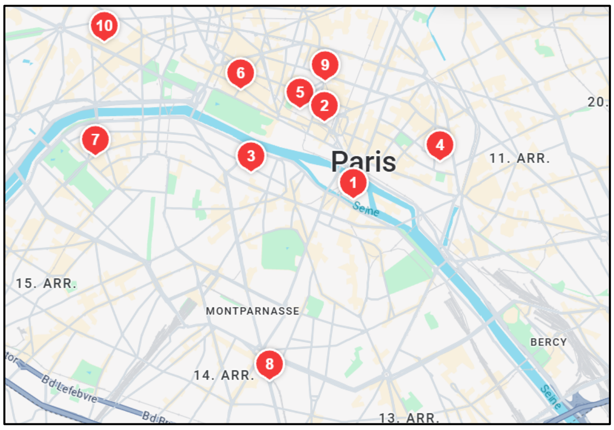

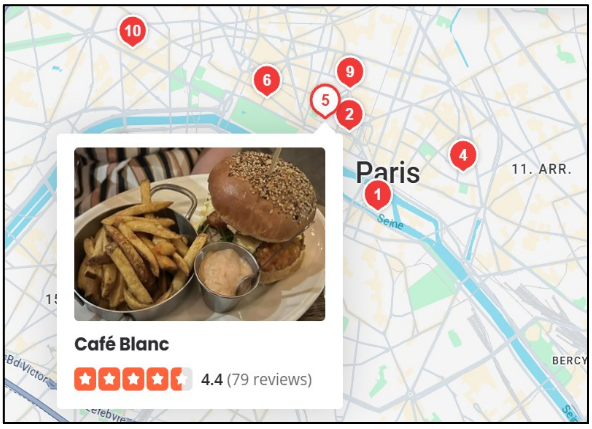

Figure 13. *Yelp map interface of a café-search in Paris, showing numbered pins (left).*
Figure 14. *The same Yelp map with hover preview (right).*

Google Maps also blends paid and unpaid visibility, using different pin shapes or colours for “Promoted Pins” and “Local Search Ads”. Nonetheless, it is unlikely that normal users know the differences between all pin designs and are actually able to recognize them as advertisements. Within the real-estate category, ImmoScout24, Rightmove.co.uk, and Zillow all sell premium or featured placements that ensure paid listings show up first before standard results. ImmoScout24.at at least explicitly discloses in its separate terms page that sponsored listings are prioritised, however the visual distinction also remains minimal in the map. On some of these platforms, such as Rightmove.co.uk, advertisements are visually marked by colour in the list view but not in the map view, which is why it was still coded as Camouflaged Advertising. In nearly all cases, the explanation of how advertising occurs exists primarily in policy or support pages, rather than at the point of interaction – the search interface – meaning that for users, these placements are effectively indistinguishable from natural map results. On Agoda, sponsored listings are technically marked, yet the labels are barely recognizable in the interface. Two exceptions stand out: Booking.com, which maintains clear labels for promoted results (“Ad” or “Sponsored”) across both list and map views, and Airbnb, which does not employ paid advertising at all. User reviews and star ratings were found across nearly all map interfaces that include user feedback elements. While it is impossible to verify if a review is real or not, the likelihood that the deceptive pattern Fake Social Proof exists can be inferred by certain questionable interface decisions. Some platforms, such as Agoda, Booking.com, and Expedia, display average scores or short popularity labels (e.g., “Highly rated,” “Popular choice”) directly on map previews or pins, but without revealing how many reviews these ratings are based on. Similarly, Tripadvisor and Yelp show star icons or ratings in extra bubbles on map markers and side panels, which gives a strong impression of reliability even though the sample size and review recency etc. is hidden from view. Google Maps has the largest user base and most well-known rating system, subsequently those reviews are also displayed. However, it goes a step further by displaying the first sentence of a review in addition to other general information on the map, but the chosen review is neither the newest nor the most highly rated, and no criteria for its selection are visible, which makes it opaque and potentially misleading. By contrast, Bing Maps uses third-party reviews, but lists its review sources transparently (Facebook, Yelp, Tripadvisor). Lastly, the real-estate platforms Rightmove.co.uk, Zillow, and Sotheby’s International Realty omit user feedback entirely and were therefore not coded for this pattern.

Overall, the differences between platform categories show that the implementation of possibly deceptive design depends less on the general business model and more on how monetization is built into the map itself. Real-estate platforms, for instance, often use default sorting that gives preference to paid or agency-verified listings, while POI platforms seem to rely on algorithms that boost certain places based on review visibility and engagement. Across these very different services, however, one thing remains uniting: ambiguous design has become a 'normal' part of how map interfaces work.

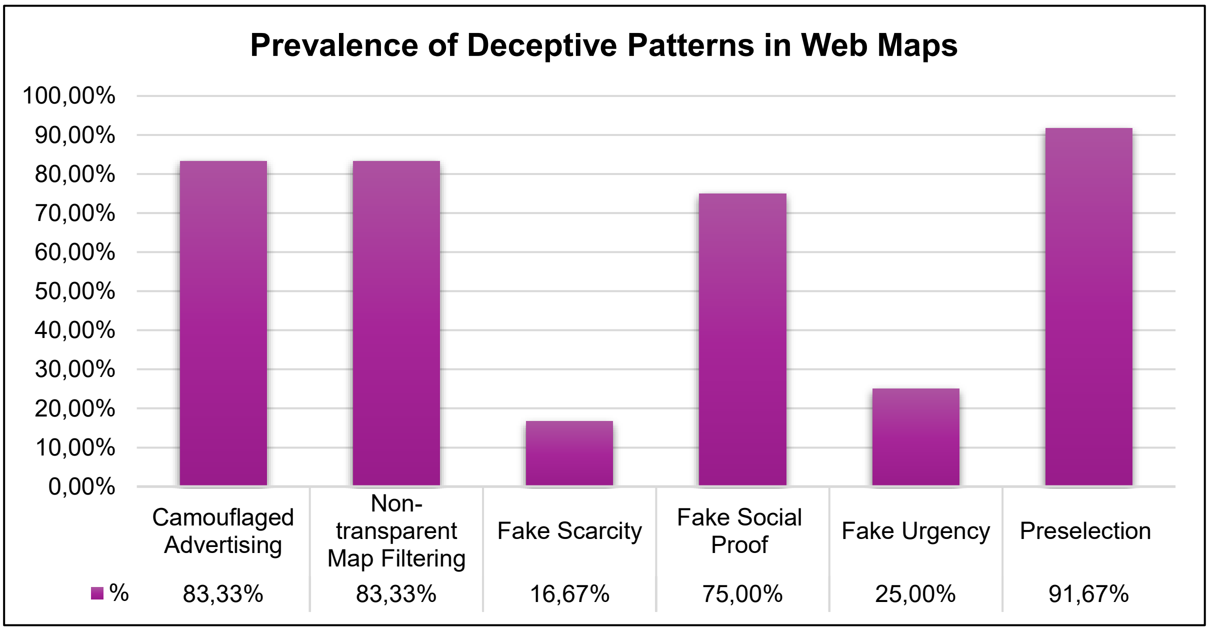

Figure 15. *Prevalence of Deceptive Patterns in Web Maps.*

The prevalence of each pattern is shown in *Figure 15*. The results confirm that Preselection, Non-transparent Map Filtering, and Camouflaged Advertising are by far the most frequent forms of deception in map-based interfaces, each appearing on more than 80 % of the analysed platforms. Dubious displays of reviews (in connection with Fake Social Proof) also occur widely (75 %), reflecting the strong influence of review-based visibility. In contrast, Fake Urgency and Fake Scarcity remain comparatively rare, which is partly because these patterns rely on real-time inventory or booking data, which only a few map-based platforms manage directly (mostly hotels & accommodations). Most POI or real-estate platforms do not handle live availability, so they cannot credibly display such tactics.

To compare the overall prevalence of deceptive patterns across platforms, a Platform Deception Risk Index was calculated for each platform *(Figure 16)*. It represents the percentage of coded pattern dimensions on which a platform scored 1 (present), and rather than indicating deliberate manipulation, reflects how frequently interface features that may reduce transparency or shape commercial visibility appear within a single platform.

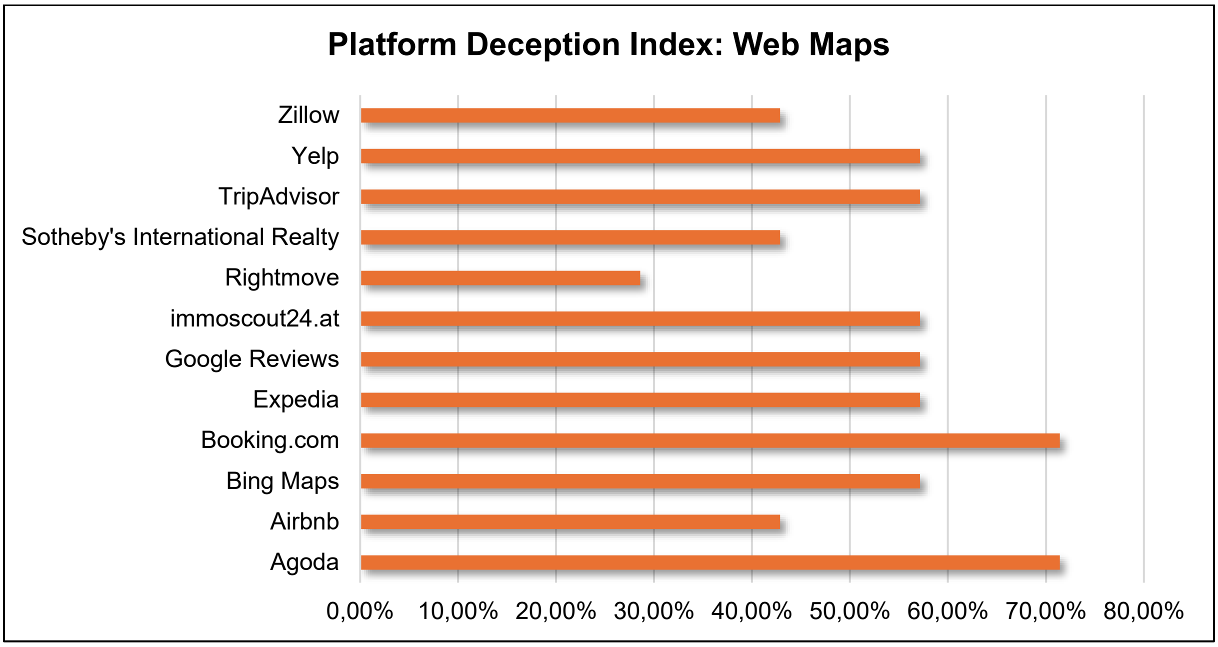

Figure 16. *Web Maps Platform Deception Index.*

The results show that deceptive design is not evenly distributed. Agoda and Booking.com reach the highest overall scores, where deceptive elements are present on roughly three quarters of the coded dimensions. TripAdvisor, Yelp, and Google Reviews follow closely, which are all driven mainly by the integration of social proof features and paid visibility listings. This may make it considerably more difficult to understand whether results are driven by relevance or commercial considerations. By contrast, Rightmove and Sotheby's International Realty rank lowest, which is due to their more information-oriented design and limited use of advertising or ranking personalization. Airbnb also scores relatively low due to the absence of paid advertising and scarcity mechanisms in its interface.

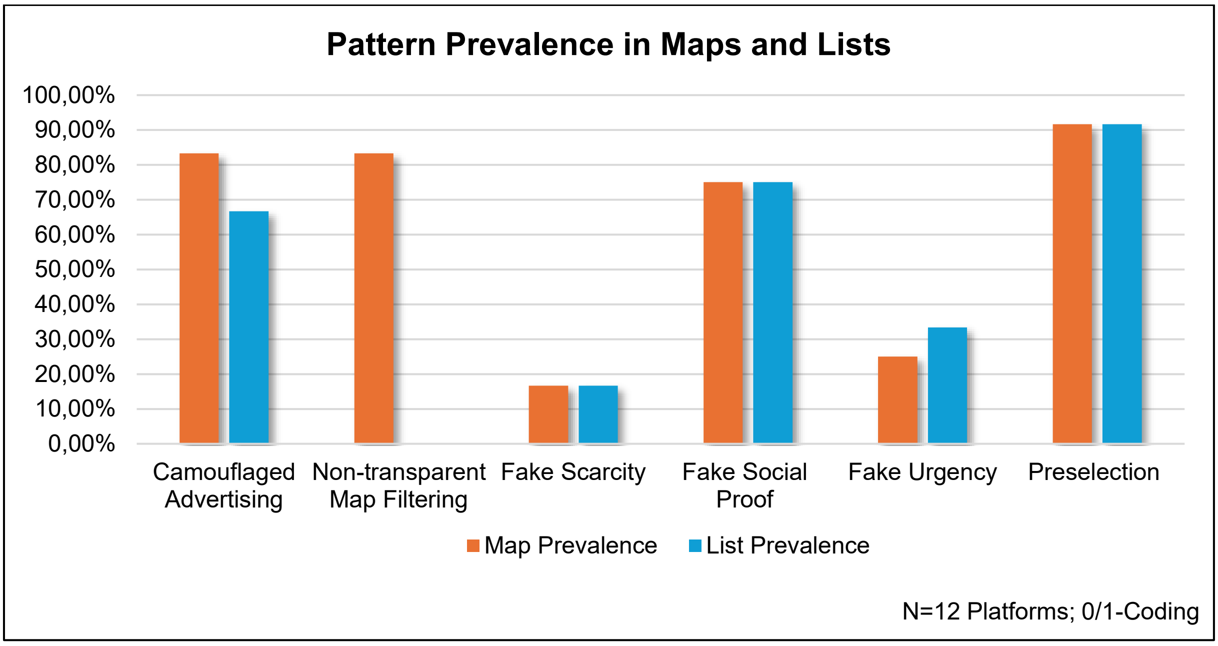

Figure 17. *Pattern Prevalence in Maps and Lists.*

Although this study focuses primarily on deceptive patterns in map-based search, the platforms analysed also present users with list-based representations of the same content. To contextualise how deceptive interface strategies show up across these types of interfaces, a comparison was conducted for the subset of patterns that can appear in both maps and lists. However, many deceptive design techniques commonly found in lists – such as oversized highlight frames that draw attention to promoted options, urgency patterns like persistent "Only 2 left!" banners, or misdirection patterns like visually emphasised price anchors (Brignull 2024) – cannot be meaningfully displayed on maps without causing readability problems. Subsequently, they were excluded from the cross-modal comparison, which focuses only on the patterns that can realistically occur in both modalities. Nonetheless, including this comparison is valuable, as it shows how interface constraints shape the set of deceptive strategies available to platforms. The results show clear structural differences: map views exhibit particularly high levels of Preselection and Non-transparent Filtering, whereas lists display stronger expressions of urgency cues and enhanced social proof. The purpose of this comparison was not to quantify which interface is potentially "more deceptive," but to illustrate how interface constraints shape the available repertoire of deceptive design strategies. The finding that lists support a broader range of persuasive and potentially deceptive techniques than maps may also help explain why several platforms have, in recent years, shifted back toward list-dominant interfaces: lists simply provide more opportunities to implement dark patterns.

4.2.1. Category-specific results: (Analysis)

A closer look at the three platform categories, Hotels & Accommodations, POIs, and Real Estate, reveals that dark patterns are not uniformly spread out, but adapted to each domain's economic goals. The following figures illustrate how the balance between manipulation and information visibility shifts depending on whether the platform prioritizes booking, engagement, or listings.

Figure 18. *Analysis Results of deceptive patterns in hotels & accommodations.*

In the Hotels & Accommodations category, Preselection, Non-transparent Map Filtering, and Fake Social Proof appear on all analysed platforms. The dominance of these patterns reflects how strongly the industry relies on algorithms and competition for bookings. Camouflaged Advertising occurs less consistently but still shows that promotional content is often blended with organic search results. This is the case for both list-view results and map-view results; however, in the map view, visual labelling clues are even weaker. Fake Scarcity and Fake Urgency occur too, to accelerate booking decisions with time-limited offers and low-stock alerts. Their presence can be explained by the transactional nature of the sector, where real-time availability allows platforms to integrate live data into persuasive messaging. Out of all categories, this one was the most susceptible to dark patterns as all of them occur.

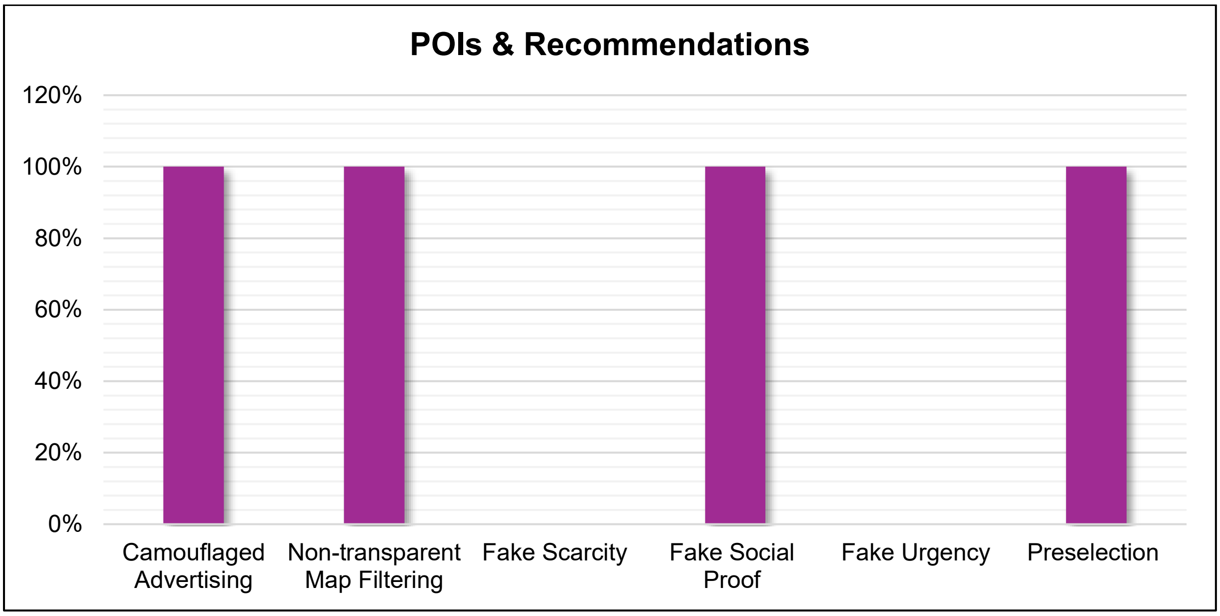

Figure 19. *Analysis Results of deceptive patterns in POI & recommendation platforms.*

Among POI & Recommendation platforms, deceptive mechanisms focus more on visibility rather than persuasion. Camouflaged Advertising, Non-transparent Map Filtering, Fake Social Proof, and Preselection are consistently present, showing how paid or algorithmically favoured locations dominate search exposure. The absence of Fake Scarcity and Fake Urgency matches the informational purpose of these services: users only explore, rather than engage in transactions. Since POI maps rely on engagement metrics, reviews, and sponsored visibility to make money, manipulation tends to occur through ranking logic and social validation through reviews rather than direct sales pressure.

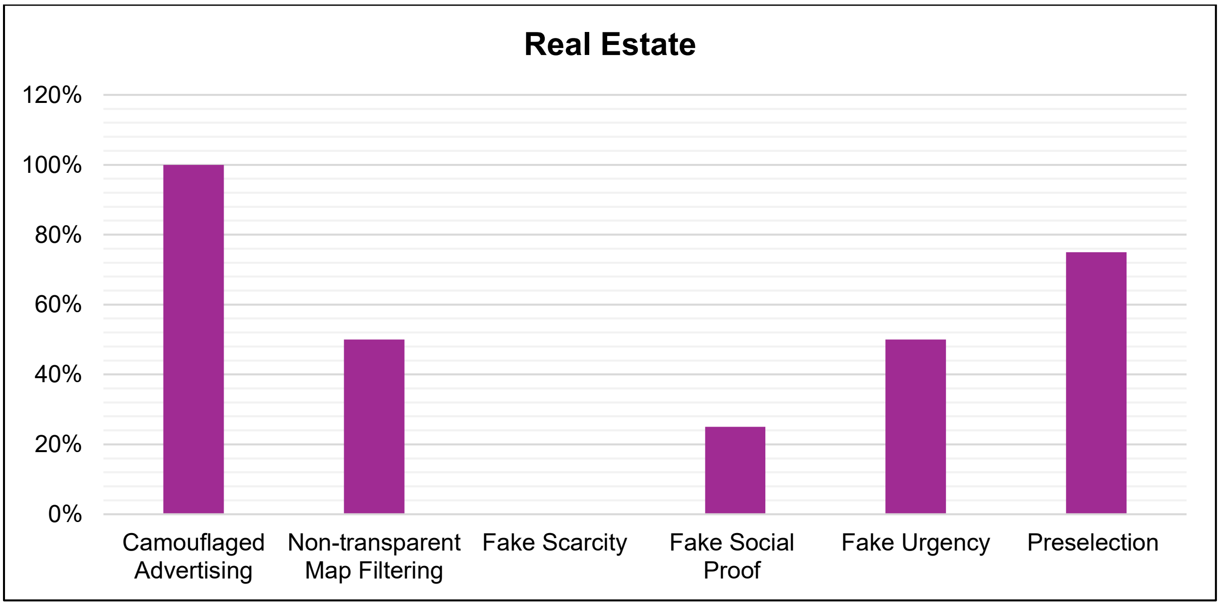

Figure 20. *Analysis Results of deceptive patterns in real estate.*

Deceptive design in *Real Estate* maps takes a more structural form. Camouflaged Advertising and Preselection are the most common patterns, using paid placements and premium listings as their main tool to get users to interact with their platforms. This ties into this category's monetization model, where visibility is sold to agencies rather than driven by user engagement. Fake Urgency appears occasionally, for example through "available" or "just listed" markers, which create a sense of market scarcity without explicitly misleading users. The rarity of Fake Social Proof is expected, as user digital reviews play a small to non-existent role in property search. Overall, real-estate platforms rely less on persuasive messaging and more on selective filtering.

4.3 Alignment Between User Perceptions and Platform Design

To synthesise the findings from the survey and the cross-platform analysis, *Table 5* below presents a convergence-divergence overview that combines the observed presence of patterns in the platform audit with the corresponding levels of perceived use in the survey. This mixed perspective is necessary because the audit reveals which mechanisms platforms objectively implement, while the survey captures which of these mechanisms are noticed or distrusted by users. Interpreting these sources separately may result in some mismatches between design intention and user experience. Therefore, integrating them provides a more complete and empirically grounded understanding of the perceived manipulation in map-based web search.

The table classifies each pattern within each platform category according to three criteria: audit presence (high, medium, low), survey perception (strong, moderate, weak/divided), and the resulting relationship (convergence, partial convergence, divergence). Audit presence shows how frequently a pattern is used on the platforms based on the coding matrix. Survey perception is derived from agreement levels, central tendency, and dispersion (SD), following the criteria established in Chapter 4.1. Convergence is recorded when both datasets indicate a similar strength of presence or awareness; partial convergence occurs when users recognise a mechanism but not to the same intensity as the audit suggests (or vice versa); and divergence appears when platforms employ a pattern a lot, but users do not perceive it. Or, alternatively, when users perceive a pattern strongly despite limited audit evidence.

Pattern	Hotels & Accommodations	POI / Recommendations	Real Estate	Interpretation
Camouflaged Advertising	**Audit: High**	**Audit: High**	**Audit: Medium–Low** (less advertising)	Very strong perceived demand for labelling across all categories; advertising is visible and palpable even where not strongly implemented.
	Survey: Strong (M ~4.0; SD <1; ≥75% agree)	**Survey: Strong** (M >4.1; SD <0.9)	**Survey: Strong** (M = 4.44; SD = .55)	
	→ Convergence	*→ Convergence*	*→ Partial Convergence (perceived > actual)*	
Fake Social Proof (Ratings/ Reviews)	**Audit: High**	**Audit: High**	**Audit: Medium**	Clear cross-category result that transparent reviews are essential. Users expect it everywhere.
	Survey: Strong (M 4.3; SD ~0.6)	**Survey: Strong** (M 4.26; SD 0.76)	**Survey: Strong** (M 4.29; SD 0.69)	
	→ Convergence	*→ Convergence*	*→ Partial Convergence*	
Preselection (Default Sorting / Algorithmic Ranking)	**Audit: High**	**Audit: Medium–High**	**Audit: Medium**	Largest perception gap in hotels: despite heavy use, users underestimate preselection.
	Survey: Weak (M = 3.07; SD = 0.79)	**Survey: Moderate/Divided** (M = 3.82; SD = 1.13)	**Survey: Weak/Divided** (M = 3.04; SD = 1.11)	
	→ Divergence	*→ Partial Convergence*	*→ Partial Convergence*	
Non-transparent	**Audit: High**	**Audit: High**	**Audit: Medium**	Users value if all results are

Map Filtering	**Survey: Moderate** (M = 3.59; SD = 0.85)	**Survey: Moderate–Mixed** (M = 4.28 BUT SD = 1.13 → divided)	**Survey: Moderate–Mixed** (M = 4.42 but SD 0.69 for price/info; completeness item 2.91 w/ SD 1.28)	shown, but often don't notice how filtering actually works. Filtering is structurally present but perceptually inconsistent.
	→ Partial Convergence	*→ Partial Convergence*	*→ Divergence*	
Fake Scarcity / Urgency	**Audit: High/Low**	**Audit: Low**	**Audit: Low** (reserved vs. available)	Users do not consistently trust urgency cues, even when heavily implemented (hotels). Perception depends strongly on context (booking vs. everyday POI).
	Survey: Weak/Divided (M = 3.26; SD = 1.05)	**Survey: Moderate** (closed vs open POI: 3.64; SD = 1.06)	**Survey: Moderate** (M = 4.18; SD = 1.09 — still divided)	
	→ Divergence	*→ Partial Convergence*	*→ Partial Convergence*	

Table 5. *Alignment Between User Perceptions and Platform Design.*

Overall, the table shows that users reliably notice design elements that are easy to see on the screen, like review displays and advertising labels, and these features match strongly between what platforms implement and what users recognise. Patterns like Camouflaged Advertising and Social Proof are both widely implemented and widely recognised, which suggests that clarity indicators (for example, a star rating on a map pin or the price) remain highly visible to users. In contrast, patterns that are more 'hidden', such as preselection and viewport-dependent filtering, show partial to substantial divergence, especially within hotels & accommodations. Here, the studied platforms make extensive use of default rankings and filters based on availability and urgency. Users seem to notice these mechanisms only inconsistently.

Context also matters: urgency messages are met with scepticism in booking environments, while availability labels appear genuinely useful in real-estate search. Taken together, these findings highlight a basic difference between what is visible and easy to interpret, and what remains hidden.

5. Discussion

Across all results, users showed strong expectations for transparency, while attitudes toward automation, completeness, and default sorting were far more mixed. The platform audit showed widespread use of visible cues like advertising and reviews, but heavy reliance on hidden mechanisms such as preselection. It is worth noting, however, that preselection is not inherently a deceptive pattern – default settings can serve legitimate usability purposes, but become problematic when they prioritise commercial interests over user preferences without transparency. Together, these insights form the basis for the interpretive discussion that follows.

The findings suggest a consistent pattern: respondents tended to report reliance on what is immediately visible on screen, while the underlying mechanisms shaping those displays appeared to go largely unnoticed. Everyone has had the experience of searching up ‚restaurants nearby‘, and then choosing one we believe to be the most suitable option, even if in many cases we end up being disappointed. This is not surprising when considering how users typically interact with digital interfaces. Prior research suggests that users tend to scan for the most prominent pieces of information – ratings, prices, labels – rather than reading interfaces in a linear way, and may treat these as the primary signals for deciding whether they trust a result (Nielsen 1997; Norman 2013: 202). In this study, ‚trust‘ is not an explicit variable but inferred from users' reliance on visible interface cues, as reflected in survey response. This understanding of trust as implicit reliance on visible cues aligns with how trust has been conceptualised in the Web GIS domain. Skarlatidou et al. argue that the complexity of spatial interfaces and non-experts' limited knowledge of GIS operations make trust particularly consequential in this context (2011: 2) – not least because spatial data inherently carries uncertainty, meaning users must rely on the system without being able to fully verify its outputs. Choosing a result without consulting additional explanations or filters can thus be understood as an implicit form of trust in the interface. What happens behind the interface, in terms of ranking or automated filtering, is far harder for users to pick up on, even though these processes may shape outcomes considerably. Background mechanisms leave no perceptible trace, meaning trust may operate less as a deliberate evaluation and more as an implicit reliance on what is immediately visible.

A similar pattern has been observed in search behaviour. A 2014 eye-tracking study found that 17% of users looking for specific information on a search engine examined only the first result and did not look anywhere else on the page (Pernice, Whitenton, & Nielsen 2014: 137). Additionally, Tversky argues that spatial representations invite direct 'reading off' of structural information such as location and distance, without requiring deeper inference (2004: 227), which may contribute to a tendency to perceive map results as more complete or neutral than they are. The survey findings are consistent with this: respondents generally expected maps to display all available results, which may reflect an assumption that what is shown corresponds to everything relevant. The platform audit suggests that this assumption is frequently unwarranted.

In contrast, the mechanisms most critical for shaping visibility on map-based platforms – ranking defaults, filtering, preselection, zoom-dependent visibility etc. – are algorithmic rather than specific UI design choices. While this is generally a broader problem of algorithms, that users see the outcomes of algorithmic decisions but are not aware that a mechanism is shaping them (Burrell 2016: 2), it also aligns with what Eslami et al. describe as the problem of “invisible algorithms”: systems that influence user experience profoundly while leaving no perceptible trace (2015: 153). From a GIS Science perspective, this invisibility is further reinforced by how algorithmic processes are embedded within spatial representations. Xing & Sieber argue that in GeoAI, computational decisions often become blurred precisely because they are presented through geographic outputs that appear neutral, making the role of algorithms difficult to perceive even when they significantly shape outcomes (2023: 627). They further say that explainability is dependent on the audience, meaning that even when explanatory mechanisms exist, they may remain meaningless to non-expert users interacting with spatial interfaces (628). The challenge gets amplified in cartographic interfaces. Traditionally, maps have long enjoyed a reputation for accuracy and have been viewed as “trusted sources of information” (Griffin 2020: 7), although no map can ever tell the whole truth. Haklay similarly argues that because maps carry an inherent sense of authority, people tend to assume what appears on a map corresponds directly to the physical world (2010: 2). When applied to map-based web search, these factors may help explain why survey respondents appeared to accept visible information such as prices and labels, as straightforward facts. This tendency aligns with findings from empirical research on trust in maps more broadly. Prestby demonstrates that users frequently rely on surface-level design cues to judge a map's trustworthiness, a process he links to the *halo effect*: when a map appears professional and visually polished, users extend that positive impression to its content as a whole, often without critically evaluating the underlying data or its source (2025: 8). In commercial map interfaces, this means that high visual quality can function as a trust signal that pre-empts scrutiny of the algorithmic processes shaping what is displayed.

Booking platforms appear to make extensive use of persuasive interface elements, and the survey results suggest that responses to these cues were notably divided in the hotel category. This dynamic is consistent with the possibility that repeated exposure may have normalised urgency cues, discount banners, or “only X rooms left” messages” for some users, while increasing scepticism in others. Even so, research suggests that repeated exposure to such techniques can make people more sceptical and lower a platform’s perceived credibility over time (Mathur et al. 2019; Fogg 2003). This may be why user trust appears relatively low in this category: people sense that they are being pushed, even if they do not fully recognise the hidden ways their choices are being shaped. Across all categories, a persistent mismatch also emerges between perceived autonomy and actual autonomy. The interactivity of map interfaces – panning, zooming, dragging – may create an impression of open-ended exploration. In practice however, the user’s sense of control is limited because the results displayed are shaped by factors users rarely have visibility into. What shows up on the map depends on things users rarely think about: how far they’ve zoomed in, how the platform decides to rank results, or which businesses have paid to appear first. Burrell calls this “opacity as intentional corporate or state secrecy” (2016: 3).

Before turning to the wider implications, it is useful to revisit the three research questions to understand how the findings collectively address them.

Research Question 1: To what extent do user expectations align with the presence of deceptive patterns in map-based web interfaces?

The findings indicate a partial and uneven alignment between user expectations and the design of map-based web interfaces. Survey results show that users generally expect map results to be transparent and complete on information. These expectations suggest that users approach map-based search as though it were an objective representation of space, rather than a curated or persuasive interface. Alignment is strongest where platforms use overt signals, such as clearly labelled advertisements. By contrast, the platform audit shows that key mechanisms shaping visibility, like default rankings or filtering logic, remain largely unnoticed by users. As a result, user expectations are met at the level of what is immediately visible, but diverge from how results are actually selected.

Research Question 2: Which potentially deceptive design patterns are commonly used on commercial map interfaces, and which of those patterns are users aware of?

The cross-platform analysis identified a clear set of deceptive patterns across commercial map interfaces, all of which suggest that, on map interfaces, potentially deceptive influence operates primarily through the regulation of visibility. Common patterns include Non-transparent Filtering, Default Ranking, Preselection, Camouflaged Advertising, and Social Proof features. Unlike list-based deception, these techniques work by shaping what enters the user's field of view rather than by using urgency or scarcity messages. Ironically, survey results suggest that users feel confident in understanding how maps work while simultaneously misjudging which elements are controlled by the system. This indicates a systematic gap between how well users believe they understand map interfaces and how much of the underlying algorithmic logic actually shapes what they see.

Research Question 3: How does the shift from spatial maps to (AI-driven) recommendation lists affect user control and transparency in location-based search?

The findings indicate that the observed shift from map-centric interfaces toward list-based and AI-driven recommendation formats may be associated with reduced user autonomy and transparency in location-based search. While maps offer users a certain degree of exploratory control through panning and zooming, lists centralise decision-making within the ranking algorithm. Lists also appear to support a broader range of deceptive design strategies, which maps cannot easily accommodate without disrupting spatial readability. Whether this structural difference contributes to platform decisions to shift away from map-based interfaces is plausible but cannot be established from the present data alone.

The platform analysis further reveals a structural shift away from map-based views toward list-dominated interfaces across a range of websites, including platforms that historically integrated spatial representations like *Willhaben, Trustpilot*, and various other review and retail platforms. This transition occurs even in contexts where geographic information would plausibly support user decision-making. A further complication is that it is increasingly difficult to determine whether list interfaces are already using AI-driven personalisation. Many platforms employ ranking systems that are adaptive, but do not disclose on what sort of mechanisms they rely on (e.g. machine learning, heuristic rules, or hybrid systems). As a result, users may have limited means to assess how items are ordered, and typically receive no information about whether rankings vary across users or context. This ambiguity further undermines transparency and limits users' ability to evaluate relevance, bias, or completeness. The transition toward list-based interfaces hence widens the gap between perceived and actual control in location-based search.

5.1. Implications for Consumer Rights

The findings of this thesis have clear implications for digital consumer rights, particularly in relation to how current regulatory frameworks address manipulative interface design. The audit identified patterns consistent with deceptive design in map-based search, while the survey suggested limited user awareness of these mechanisms. Taken together, these findings indicate that some map interfaces may very well fall within the scope of dark patterns as defined in EU consumer and platform law. Yet map-based search is not an explicitly recognised risk area in current dark-pattern legislation, which largely focuses on consent flows, subscriptions and e-commerce nudging (Yi & Li 2025: 3). This regulatory blind spot is problematic insofar that deceptive patterns in maps may shape what users see in ways that are functionally comparable to more recognised forms of dark patterns such as countdown timers or manipulative cookie banners. However, since many of the identified deceptive patterns in this thesis operate through default visibility rather than explicit interaction, they are unlikely to trigger user complaints or attract regulatory scrutiny. From an enforcement perspective, this makes map-based manipulation particularly difficult to identify, as there is often no single misleading statement or prompt to evaluate. Instead, influence emerges cumulatively through what is shown, prioritised, or omitted.

Recent regulatory developments reinforce the relevance of this concern. The Digital Services Act now extends consumer protection obligations to interface design and recommender systems. For instance, the DSA explicitly prohibits the use of deceptive patterns that materially distort or impair users' ability to make informed decisions and requires platforms to provide transparency around the main parameters of recommender systems (European Union, 2025). While these provisions mark a significant step forward, the findings of this thesis suggest that they may be difficult to apply to map-based search, where influence emerges through default spatial visibility rather than through discrete interface elements or explicit user choices. At the same time, the European Parliamentary Research Service has noted that current legal instruments remain fragmented and insufficiently adapted to complex, hybrid interfaces (EPRS 2025).

Against this backdrop, effective regulation may need to look beyond individual interface elements and consider how default visibility and ranking shape choice environments over time. For map-based search, this raises the question of whether transparency obligations should extend not only to labels and disclosures, but also to the principles governing spatial ranking in recommender systems. Regulatory approaches need to account for how such rules structure what users are able to see and choose.

5.2. Implications for GeoAI

The findings of this thesis also help explain how AI-mediated spatial search may develop in the future. As already pointed out, users strongly perceive visible features in maps but struggle to detect elements like Preselection and Non-transparent Map Filtering. This "perception gap" becomes increasingly consequential as GeoAI systems take on more of the ranking, filtering and decision-making work in map interfaces. Recent research on autonomous GIS shows that next-generation systems will be capable of independently retrieving data, generating workflows, selecting models, and optimising outputs based on internal performance metrics (Li et al., 2025: 8). In commercial settings, such optimisation tends to prioritise engagement, revenue, and sponsor visibility. When combined with the opacity documented in this thesis, AI may amplify some of the inherent structural issues that are already present in map-based search. Once future spatial interfaces rely more and more on AI components that go beyond selecting and ranking items – operating with greater degrees of autonomy – the underlying processes may become even less clear, particularly when outputs are personalised and updated in real time (Kitchin 2017; Mathur et al. 2019). This is the crux of the algorithmic auditing in general: as systems become more complex, users' ability to understand or influence design choices decreases, while the system's influence over their decision-making increases.

Building on this, GIScience scholarship suggests that 'trust' in spatial systems increasingly shifts from the map as a representational artefact to the computational processes that generate it. Writing in a pre-AI context, Sieber & Haklay say that rather than evaluating how geographic information is produced, users often delegate epistemic authority to the system as a whole, assuming that spatial outputs are both technically sound and contextually appropriate (2015: 126). They also note that in digital geographic systems, trust is frequently embedded in workflows and infrastructures rather than consciously assessed by users (2015: 129). In GeoAI contexts, this delegation of trust becomes particularly consequential as algorithmic decisions are not only hidden but continuously adapted over time.

At the same time, GeoAI research indicates that spatial decision-making will increasingly be shared between users and automated systems. Autonomous GIS frameworks anticipate agents capable of data retrieval, map generation, pattern detection, and iterative model refinement (Li et al. 2025: 508). Many of these techniques already underpin recommendation systems and routing. As they become more integrated into map interfaces, the boundary between "map" and "list" is likely to blur more and more, with AI systems pre-structuring what is shown, in what order, and under which conditions. And if visibility decisions stem solely from AI, traditional "disclosure-style-transparency", which this study has shown to be insufficient in practice, will not help users understand how choices are shaped. This trend is not merely theoretical. It is consistent with the study's findings on the dominance of list-based interfaces across the analysed platforms. In addition, while trying to select the platforms for this study, it became apparent that several large and widely used services no longer offer an interactive map interface at all – among them Trustpilot and other platforms – despite operating clearly in geographic domains. Instead, on many websites, results are presented exclusively through ranked lists, only in some cases are they supported by a minimal or static map thumbnail. From a commercial perspective this makes naturally sense: lists are way easier to monetise. The observed decline of maps is therefore less an innocent design update than a structural move away from spatial transparency. It is plausible that as AI-driven ranking becomes more prevalent, the shift toward list-based interfaces may continue, though the pace and extent of this transition will depend on factors beyond the scope of this study. Recent work on geographic foundation models and large-scale GeoAI systems further underscores these concerns. Increasingly, spatial representations are produced by models trained on heterogeneous data sources (Xing & Sieber 2023: 638). These sources differ widely in quality and coverage, so variation affects how geographic patterns are represented. Ethical discussions in GeoAI research and Spatial Data Science emphasise that such systems do not merely reflect geographic reality but take part in shaping it. AI systems can both motivate participation in spatial decision-making and act as the medium through which spatial information is presented, thereby privileging certain places, actors, and relationships (Sieber et al. 2024: 1390). Additionally, AI’s complex computational processes are often shielded from scrutiny, and users are implicitly encouraged to defer judgment to technical expertise rather than engage critically with how outcomes are produced (2024: 1395).

On these grounds, transparency standards will need to address algorithmic behaviour directly: how and why first results are displayed, which variables guide ranking, and how commercial incentives influence spatial outputs. Calls for responsible AI in GIS emphasise verifiability, interpretability and human oversight (Li et al. 2025: 529), all of which become essential as AI takes a more active role in structuring map-based search. Overall, the findings in this thesis offer a small but useful empirical basis for future AI research by showing where users tend to overlook structural interface mechanisms, and thus where automated spatial ranking may require closer scrutiny.

6. Limitations and Outlook

Like any study dealing with commercial digital platforms, this research comes with a number of unavoidable limitations. The platform analysis captures interfaces at specific moments in time and cannot fully account for continuous updates, A/B testing, where platforms test different interface versions, or ongoing personalisation of content – all of which may subtly change what users see. This limitation could have been mitigated through repeated observations over longer periods or by systematically documenting interface changes across multiple sessions and user profiles; however, such an approach would have exceeded the available time and resources. Similarly, the internal logic behind ranking and filtering systems remains inaccessible, meaning that the analysis is limited to observable interface behaviour. The survey component also relies on self-reported perceptions, which may be well-suited to identifying expectation gaps, but may not always reflect how users behave in real-world decision situations, especially under time pressure or commercial influence. This also links to the caveat that survey responses may have differed had participants been shown actual map examples rather than abstract statements, as they may not have fully grasped the practical implications of certain design choices. The audit was also conducted by a single researcher, which is typical for exploratory studies of this scope. While the coding criteria were defined in advance to ensure consistency, intercoder reliability testing was not feasible within the constraints of this project. Future studies could strengthen this aspect by employing multiple coders and calculating reliability metrics such as Cohen's kappa. In addition, the sample size was small due to time and resource limitations. Accordingly, the results cannot be easily generalised beyond the platforms and user groups included in the study. A larger and more diverse sample could have reduced this limitation and allowed for more fine-grained subgroup analyses, for example across age groups or levels of digital literacy. Lastly, the analysis focused on a selected set of platforms and on the dark patterns identified in this study. As this is one of the first systematic attempts to examine deceptive patterns specifically in map-based search, the classification developed here is necessarily exploratory. It does not claim to capture the full diversity of spatial search interfaces, nor to exhaust the range of manipulative strategies that may exist in cartographic or location-based systems. These constraints are less a flaw of the study than a reminder of how difficult it is to research continuously changing digital environments.

A conceptual limitation worth noting separately is the question of intent. While this study identifies interface features that may reduce transparency or shape user perception in commercially convenient ways, it cannot establish whether these patterns result from deliberate design decisions or from technical constraints, standard industry practices, or unintended side effects.

Despite these limitations, the study opens several directions for future research. One promising path is to examine how people actually use map-based search in everyday situations, for example through behavioural experiments, longitudinal observations, or eye-tracking studies. Eye-tracking has proven especially valuable for studying visual attention and cognitive load in spatial interfaces (Kiefer et al. 2017: 7) and could reveal how users scan and respond to ranking, filtering, and preselection mechanisms under conditions of time pressure or limited attention. Such observational approaches were initially planned as a third empirical component of this project, but were ultimately set aside to keep the scope manageable for a master's thesis. However, future studies could build on this foundation by combining the survey data with task-based experiments, clickstream analysis, or controlled search scenarios. As map interfaces increasingly rely on adaptive and AI-driven systems, future work could also explore how personalisation and automated recommendations influence visibility and perceived control over time. In particular, future research could examine how repeated exposure to personalised rankings influences how much users trust map results and how strongly they come to rely on recommendations instead of actively exploring the map. This is especially relevant for GeoAI-based systems, as their models continuously adapt to user behaviour and may gradually narrow the range of locations that users are shown. In addition, further studies could explore whether users are aware that AI systems shape these results at all. Another useful direction would be comparative research across legal contexts. Examining how different regulatory frameworks address manipulative design in maps and whether certain protection measures improve transparency in practice is an essential part of ensuring that consumer protection keeps pace with how digital maps increasingly shape everyday choices. Comparative studies could examine how map platforms behave under different regulatory regimes, such as in the European Union, the United States, or other regions, and whether stricter rules actually lead to more transparent map interfaces. Rather than focusing only on legal texts, such research could look at what users really see on their screens and test whether regulation makes manipulative design harder to hide or simply encourages platforms to disguise it more carefully. Beyond that, there are currently many non-traditional forms of spatial interaction emerging, such as voice-based assistants, AI-generated summaries, and augmented reality navigation systems. They raise new questions about visibility, explainability, and user agency. Investigating whether deceptive or manipulative patterns persist, change, or intensify in these interfaces would be a valuable extension of this work. Finally, future research could further develop the theoretical integration between critical GIS, platform studies, and dark pattern research. By combining these perspectives, researchers may better understand how spatial data, interface design, and commercial incentives interact and affect us.

7. Conclusion

This thesis set out to examine a part of digital life that is used daily, trusted almost instinctively, and questioned surprisingly rarely: map-based web search. While maps are often perceived as neutral representations of space, the findings of this study demonstrate that contemporary digital maps are not neutral.

To address this gap in how deceptive design in map-based search is conceptualised and studied, a map-specific typology of deceptive patterns was first developed. This typology adapts and extends existing dark pattern frameworks from e-commerce and interface design to the interface-specific mechanisms and ranking logics underlying spatial search. As there has been little prior work examining deceptive patterns in maps specifically, this step represents a novel contribution. It provides a systematic basis for analysing how manipulation operates within cartographic interfaces. Based on that, this work addressed three core questions: (1) Whether users' expectations of map-based search correspond to how commercial map interfaces function, (2) which forms of deceptive design are embedded in these systems, and (3) how the growing shift toward list-based and AI-driven recommendations affects user control. The findings show that while users generally expect maps to be spatially complete, these expectations align only with what is visually explicit. The most influential mechanisms, which include Preselection and Non-transparent Map Filtering, remain largely unnoticed. Across platforms, deceptive patterns were found to be widespread, yet they differ from classic e-commerce tactics by operating primarily through selective visibility rather than open pressure or urgency. At the same time, the observed shift toward ranked lists and AI-driven recommendations may further centralise decision-making and limit users' ability to understand or influence how results are generated. Together, these results suggest that map-based search offers a strong *illusion* of control, but discreetly structures choices in ways that are neither fully visible nor easily questioned. Moreover, as GeoAI systems take on a more active role in structuring spatial search, this tension is likely to intensify rather than disappear. If AI systems take on a more active role in deciding which locations are shown, prioritised, or omitted, transparency cannot remain a matter of small labels or hidden information icons alone. Altogether, the study provides clear insights into how users perceive and interact with commercial map interfaces, but they also point to the limits of user-facing analysis, as the underlying algorithmic and organisational processes remain largely inaccessible. Ultimately, this thesis does not argue that digital maps are inherently deceptive on purpose – but the findings are consistent with the conclusion that they are quietly persuasive. Recognising this is a necessary first step toward more transparent spatial interfaces. Or, put more simply: the next time a map claims to show “the best option nearby,” it may be worth asking whose definition of “best” is actually at work.

8. Bibliography

Ali, D. E. (2020). Geographic Information System (GIS): Definition, Development Applications & Components. Ananda Chandra College, Jalpaigur, 1-13.

Brignull, H. (2023). *Deceptive Patterns: Exposing the Tricks Tech Companies Use to Control You*. Testimonium Ltd.

Brignull, H., Leiser, M., Santos, C., & Doshi, K. (2023). *Deceptive patterns – user interfaces designed to trick you.* Online 2023. https://www.deceptive.design/ (26.12.2024)

Bösch, C., Erb, B., Kargl, F., Kopp, H., & Pfattheicher, S. (2016). Tales from the dark side: Privacy dark strategies and privacy dark patterns. *Proceedings on Privacy Enhancing Technologies* (4), 237–254. https://doi.org/10.1515/popets-2016-0038

Booking.com. (2024). *Booking Holdings Inc.'s Digital Markets Act Compliance Report*. Public Summary. https://www.bookingholdings.com/wp-content/uploads/2024/11/DMA-Compliance-Report.pdf (08.05.2025)

Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society 3*(1). https://doi.org/10.1177/2053951715622512

Butler, O. (2017, December 6). I Made My Shed the Top Rated Restaurant On TripAdvisor. Vice. https://www.vice.com/en/article/i-made-my-shed-the-top-rated-restaurant-on-tripadvisor/

Chen, C., Wei, M., Zhang, P., Luo, P., & Meng, L. (2025). An infographic framework of GeoAI ethics based on news data. *Cartography and Geographic Information Science*, 1–21. https://doi.org/10.1080/15230406.2025.2593949

Chen, L., Chen, G. & Wang, F. (2015). Recommender systems based on user reviews: the state of the art. *User-Modeling and User-Adapted Interaction: Journal of Personalization Research 25*(2), 99–154. https://doi.org/10.1007/s11257-015-9155-5

Chini, M. (2024). *Only one room left: Booking.com reprimanded for misleading consumers.* The Brussel Times, 28 May 2024. https://www.brusselstimes.com/1063448/only-one-room-left-booking-com-reprimanded-for-misleading-consumers (09.05.2025)

CJEU. (2024). *Compass Banca SpA v. AGCM* (Case C-646/22).

Elwood, S. (2006). Critical Issues in Participatory GIS: Deconstructions, Reconstructions, and New Research Directions. *Transactions in GIS 10*(5), 693–708. https://doi.org/10.1111/j.1467-9671.2006.01023.x

Eslami, M., Rickman, A., Vaccaro, K., Aleyasen, A., Vuong, A., Karahalios, K., Hamilton, K., & Sandvig, C. (2015). “I always assumed that I wasn’t really that close to [her]”: Reasoning about invisible algorithms in news feeds. *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 153–162. https://doi.org/10.1145/2702123.2702556

European Union. (2022). *Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services and amending Directive 2000/31/EC (Digital Services Act)*. OJ L 277, 27.10.2022, p. 1–102. Preamble 61–70.

European Commission. (2024). Commission sends preliminary findings to X for breach of the Digital Services Act. https://ec.europa.eu/commission/presscorner/detail/en/ip_24_3761 (10.10.2025)

European Parliament & Council of the European Union. (2005). *Directive 2005/29/EC of the European Parliament and of the Council of 11 May 2005 concerning unfair business-to-consumer commercial practices in the internal market* (Unfair Commercial Practices Directive). *Official Journal of the European Union*, L149, 22–39. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32005L0029 (10.10.2025)

European Parliamentary Research Service. (2025). Regulating dark patterns in the EU: Towards digital fairness. At-a-Glance EPRS_ATA(2025)767191. European Parliament. https://www.europarl.europa.eu/thinktank/en/document/EPRS_ATA(2025)767191 (05.01.2026)

Fairpatterns.com. *“Dark Patterns” Targeted by EU Institutions: A New Era of Digital Fairness*. Online, November 18, 2024. https://www.fairpatterns.com/post/dark-patterns-targeted-by-eu-institutions-a-new-era-of-digital-fairness (16.05.2025)

Fogg, B. J. (2003). *Persuasive Technology: Using Computers to Change What We Think and Do*. Morgan Kaufmann.

Gao, S. (2021). *Geospatial artificial intelligence (GeoAI*). Oxford University Press.

Gao, S., Hu, Y., & Li, W. (Eds.). (2023). *Handbook of Geospatial Artificial Intelligence* (1st ed.). CRC Press.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology, 62*, 451–482. https://doi.org/10.1146/annurev-psych-120709-145346

Google Cloud. (2021, September 25). *Recommendations AI modeling*. Google Cloud Blog. https://cloud.google.com/blog/topics/developers-practitioners/recommendations-ai-modeling (01.02.2026)

Google Maps Platforms. (2025). Visualizing data with BigQuery and Datasets API. https://developers.google.com/maps/architecture/bigquery-datasets-visualization (03.03.2025)

Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., & Toombs, A. L. (2018). The dark (patterns) side of UX design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, Article 534, 1–14. https://doi.org/10.1145/3173574.3174108

Greenberg, S., Boring, S., Vermeulen, J, & Dostal, J. (2014). Dark Patterns in Proxemic Interactions: A Critical Perspective. *Proceedings of the 2014 Conference on Designing Interactive Systems (DIS '14)*. Association for Computing Machinery, New York, NY, USA, 523–532. https://doi.org/10.1145/2598510.2598541

Griffin, A. L. (2020). Trustworthy maps. *Journal of Spatial Information Science, 20*, 5–19. https://doi.org/10.3998/josis.2020.20.638

Haklay, M. (2010). How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets. *Environment and Planning B: Planning and Design 37*(4), 682–703. https://doi.org/10.1068/b35097

Harley, J.B. (1989). Deconstructing the Map. *Cartographica 26*(2), 1–20. https://doi.org/10.3138/E635-7827-1757-9T53

Jahagir, R. (2025). Understanding the EU's Digital Services Act Enforcement Against X. Online, Apr 5, 2025. https://www.techpolicy.press/understanding-the-eus-digital-services-act-enforcement-against-x/?utm_source=chatgpt.com (08.05.2025)

Janowicz, K., Gao, S., McKenzie, G., Hu, Y., & Bhaduri, B. (2020). GeoAI: Spatially Explicit Artificial Intelligence Techniques for Geographic Knowledge Discovery and Beyond. International Journal of Geographical Information Science *34*(4), 625–636. https://doi.org/10.1080/13658816.2019.1684500

Janowicz, K., Shi, M., Liu, Z., Beluli, N., Bischoff, G., Huber, L., Huber, P., Imrecke, M., Kastner, M., Knabe, D., König, F., Neuhauser, R., Neustätter, N., Straßburger, S., Thurnhofer, S., Venier, F., & Welz, L. (2024). *Quo vadis – Map-based Search?* Zenodo. https://doi.org/10.5281/zenodo.14260493

Kang, Y., Gao, S. & Roth, R. (2024). Artificial intelligence studies in cartography: a review and synthesis of methods, applications, and ethics. *Cartography and Geographic Information Science 51*(4), 1-99. https://doi.org/10.48550/arXiv.2312.07901

Kenney, M., Zysman, J. (2020). The platform economy: restructuring the space of capitalist accumulation. Cambridge Journal of Regions, Economy and Society, *13*(1), 55–76. https://doi.org/10.1093/cjres/rsaa001

Kiefer, P., Giannopoulos, I., Raubal, M., & Duchowski, A. (2017). Eye tracking for spatial research: Cognition, computation, challenges. *Spatial Cognition & Computation 17*(1–2), 1–19. https://doi.org/10.1080/13875868.2016.1254634

Kitchin, R. (2017). Thinking critically about and researching algorithms. *Information, Communication & Society, 20*(1), 14–29. https://doi.org/10.1080/1369118X.2016.1154087

Lewis, C. (2014). *Irresistible Apps: Motivational Design Patterns for Apps, Games, and Web-based Communities*. Berkeley: Apress.

Li, Z., Ning, H., Gao, S., Janowicz, K., Li, W., Arundel, S. T., Hodgson, M. E. (2025). GIScience in the era of Artificial Intelligence: a research agenda towards Autonomous GIS. *Annals of GIS 31*(4), 501–536. https://doi.org/10.1080/19475683.2025.2552161

Lomas, N. (2024, March 14). AliExpress is first online marketplace to face DSA investigation by EU. *TechCrunch*. https://techcrunch.com/2024/03/14/aliexpress-dsa-probe/ (09.06.2025)

Mai, G., Huang, W., Sun, J., Song, S., Mishra, D., Liu, N., Gao, S., Liu, T., Cong, G., Hu, Y., Cundy, C., Li, Z., Zhu, R., and Lao, N. (2024). On the Opportunities and Challenges of Foundation Models for GeoAI (Vision Paper). ACM Trans. *Spatial Algorithms Systems* 10(2), Article 11 (June 2024), 46 pages. https://doi.org/10.1145/3653070

Marquez, J. O., Meirelles, P., and da Silva, T. S.: Interactive Web Maps: Usability Heuristics Proposal, Proc. Int. Cartogr. Assoc., 4, 70, https://doi.org/10.5194/ica-proc-4-70-2021, 2021.

Mathur, A., Acar, G., Friedman, M. G., Lucherini, E., Mayer, J., Chetty, M., & Narayanan, A. (2019). Dark patterns at scale: Findings from a crawl of 11K shopping websites. *Proceedings of the ACM on Human-Computer Interaction, 3(CSCW)*, 1–32. https://doi.org/10.1145/3359183

Medyńska-Gulij, B., Wielebski, Ł., Halik, Ł., & Smaczyński, M. (2020). Complexity Level of People Gathering Presentation on an Animated Map – Objective Effectiveness Versus Expert Opinion. *ISPRS International Journal of Geo-Information 9*(2), 117. https://doi.org/10.3390/ijgi9020117

Medynska-Gulij, B., Forrest, D., Cybulski, P. (2021). Modern Cartographic Forms of Expression: The Renaissance of Multimedia Cartography*. ISPRS International Journal of Geoinformation 10*, 484. https://doi.org/10.3390/ijgi10070484

Muehlenhaus, I. (2014). *Web Cartography: Map Design for Interactive and Mobile Devices*. Boca Raton, FL: CRC Press.

Norman, D. A. (2013). *The Design of Everyday Things: Revised and Expanded Edition*. New York: Basic Books

Nielsen, J. (1997, October 1). How users read on the web. *Nielsen Norman Group*. https://www.nngroup.com/articles/how-users-read-on-the-web/ (15.11.2025)

OpenStreetMap Engineering Working Group. (2024). *A progress update on vector tiles from the Engineering Working Group*. Online 2023 July 6. https://blog.openstreetmap.org/2024/06/06/an-progress-update-on-vector-tiles-from-the-engineering-working-group/ (25.12.2024)

Pavlovskaya, M. (2018). Critical GIS as a tool for social transformation. *Canadian Geographer* 62, 40-54. https://doi.org/10.1111/cag.12438

Pernice, K., Whitenton, K., & Nielsen, J. (2014). *How people read online: The eyetracking evidence* (Updated edition). Nielsen Norman Group.

Peterson, M. P. (2014). *Mapping in the Cloud*. New York: The Guilford Press.

Polaris Market Research. (2024*). Digital Map Market Share, Size, Trends, Industry Analysis Report, By End-Use; By Technology; By Application; By Region; Segment Forecast, 2024 – 2032*. Online 2024. https://www.polarismarketresearch.com/industry-analysis/digital-map-market (30.12.2024)

Prestby, T. J. (2025). Trust in maps: what we know and what we need to know. *Cartography and Geographic Information Science 52*(1), 1–18. https://doi.org/10.1080/15230406.2023.2281306

Ravi, L., & Vairavasundaram, S. (2016). A Collaborative Location Based Travel Recommendation System through Enhanced Rating Prediction for the Group of Users. *Computational intelligence and neuroscience*, 1291358. https://doi.org/10.1155/2016/1291358

Sheppard, E. (2005). Knowledge production through critical GIS: Genealogy and prospects. *Cartographica 40*(4), 5–21. https://doi.org/10.3138/GH27-1847-QP71-7TP7

Sieber, R.E., Brandusescu, A., Sangiambut, S., & Adu-Daako, A. (2024). What is civic participation in artificial intelligence*? Environment and Planning B: Urban Analytics and City Scienc*e *52*, 1388 - 1406. https://doi.org/10.1177/23998083241296200

Sieber, R. E., & Haklay, M. (2015). The epistemology(s) of volunteered geographic information: a critique. *Geo: Geography and Environment 2*(2), 122–136, doi: 10.1002/geo2.10.

Skarlatidou, A., Wardlaw, J., Haklay, M. & Cheng, T. (2011). Understanding the influence of specific Web GIS attributes in the formation of non-experts' trust perceptions. *Advances in Cartography and GIScience 1*, 219–238. https://doi.org/10.1007/978-3-642-19143-5_13

Thatcher, J., Bergmann, L., Ricker, B., Rose-Redwood, R., O'Sullivan, D. et al. (2016). Revisiting Critical GIS. *Environment and Planning A 48*(5), 815–824. https://doi.org/10.1177/0308518X15622208

The Nielsen Company. (2015). Global Trust in Advertising Survey. https://www.nielsen.com/content/dam/nielsenglobal/apac/docs/reports/2015/nielsen-global-trust-in-advertising-report-september-2015.pdf (08.05.2025)

Trump, D.J. (2025). Restoring Names That Honor American Greatness. Executive Order 14172. The White House. https://www.whitehouse.gov/presidential-actions/2025/01/restoring-names-that-honor-american-greatness/ (15.04.2026)

Tversky, B. (2004). Visuospatial reasoning. In K. Holyoak & R. Morrison (Eds.), *The Cambridge Handbook of Thinking and Reasoning* (pp. 209–241). Cambridge: Cambridge University Press.

Wu, S., Henggeler, K., Chen, Y. & Hurni, L. (2024). A roadmap for generative mapping: unlocking the power of generative AI for map-making. Institute of Cartography and Geoinformation, ETH Zurich Switzerland. https://doi.org/10.48550/arXiv.2410.15770

Xing, J., & Sieber, R. (2023). The challenges of integrating explainable artificial intelligence into GeoAI. *Transactions in GIS 27*, 626–645. https://doi.org/10.1111/tgis.13045

Yi, W., Li, Z. (2025). Mapping the scholarship of the regulation of dark patterns: A systematic review of concepts, regulatory paradigms, and solutions from law and HCI perspectives. *Computer Law & Security Review 59*, 106225, 1-23. https://doi.org/10.1016/j.clsr.2025.106225.

Young, S. W. H. (2014*).* Improving Library User Experience with A/B Testing: Principles and Process. *Weave: Journal of Library User Experience 1*(1), 1-24. doi:10.3998/weave.12535642.0001.101. hdl:2027/spo.12535642.0001.101

Zhang, Q., Lu, J. & Jin, Y. (2021). Artificial intelligence in recommender systems. *Complex Intell. Syst. 7*, 439–457. https://doi.org/10.1007/s40747-020-00212-w

Zejdlik, J., & Vozenilek, V. (2025). Exploring Cartographic Differences in Web Map Applications: Evaluating Design, Scale, and Usability. *ISPRS International Journal of Geo-Information 14*(1), 9. https://doi.org/10.3390/ijgi14010009

Zhunis, B., Gartner, G., & Klettner, S. (2021). Cross-Cultural Differences in Map Design Perception. In P. Zamperlin, A. Cantile, & M. Milli (Eds.), *Abstracts of the ICA 3*, 1–3. https://doi.org/10.5194/ica-abs-3-329-2021