Teaching or Manipulating? On the Adoption of Bright and Deceptive Patterns by Duolingo

João Vítor Castro[1], **George Valença**[1]

[1]Departamento de Computação – Universidade Federal Rural de Pernambuco (UFRPE)

{joao.vitorc,george.valenca}@ufrpe.br

Abstract. Introduction*: In 2024, Duolingo surpassed 100 million active users, including a large number of children, raising concerns about its design practices.* ***Objective****: This study aims to critically examine Duolingo's interface through the lens of fair and deceptive design patterns, especially regarding children's rights.* ***Methodology or Steps****: We conducted a qualitative analysis of the platform's visual and interaction elements, focusing on engagement strategies.* ***Expected Results****: While Duolingo employs ethical practices like clear interfaces, we identified manipulative patterns such as excessive notifications and emotionally charged visuals. These findings contribute to discussions on ethical design in educational apps for children.*
Keywords *Deceptive Patterns, Fair Patterns, Children's Rights, Educational Apps, Design Ethics.*

1. Introduction

In the contemporary digital landscape, platforms and applications significantly shape user experiences and behaviors. For children and adolescents in Brazil, the internet has become a ubiquitous space, with 92% of the population aged 9 to 17 using it in 2022, amounting to 24.4 million individuals [Núcleo de Informação e Coordenação do Ponto BR 2024]. This intense online engagement, which includes activities such as listening to music (87%), watching videos (82%), researching for schoolwork (80%), using instant messaging (79%), and downloading apps (70%), expands both the opportunities and the risks inherent to digital environments [Núcleo de Informação e Coordenação do Ponto BR 2024].

In this context, the design of digital interfaces plays a crucial role, often influencing users' decisions in ways that may not align with their best interests [Potel-Saville e Da Rocha 2023]. Deceptive or manipulative interfaces, termed "dark patterns" or "deceptive patterns" by Harry Brignull in 2010, are widely used in digital platforms. These patterns are defined as *"tricks used on websites and apps that make you do things that you didn't intend to do, like buying or signing up for something"* [Brignull 2024]. The scientific literature identifies a common characteristic in these patterns: they modify the choice architecture presented to users, exploiting cognitive biases to steer users toward decisions unknowingly, against their preferences or best interests.

There is multidisciplinary evidence of the individual and structural harm these practices cause. The use of misleading or obstructive language is an example of a deceptive pattern [Sandhaus 2023]. Reflecting on it to promote a perspective of Plurality and Decoloniality while valuing the specificities of audiences such as children and

adolescents requires alignment with normative perspectives that are relevant to regulators and must be considered in studies on the topic (e.g., individual well-being or autonomy), which can present significant challenges [de Oliveira et al. 2024].

In response to this issue, a complementary field of research has emerged, aiming not only to identify and classify deceptive patterns but also to propose solutions. In this context, the notion of fair or bright patterns is introduced as design strategies to shift from a problem-oriented to a solution-oriented perspective. These concepts aim to provide usable, accessible, and sustainable tools to empower all stakeholders to act against manipulative design [Potel-Saville e Da Rocha 2023]. Bright patterns are presented as an ethical approach to counteract deceptive design [Sandhaus 2023].

In 2024, Duolingo surpassed 100 million active users, positioning itself as one of the largest and most influential digital language-learning platforms in the world, according to a report published by the company for investors. Its wide accessibility, gamified features, and friendly interface have made it especially popular among child users. However, despite educational opportunities, Duolingo employs a wide range of interactive and visual resources to encourage continued use of the platform [Shortt et al. 2023]. In many cases, these approaches are effective in promoting daily practice and persistence in the learning process. However, there is a fine line between motivation and manipulation. The way certain elements are applied, particularly in terms of repetition, urgency, or emotional appeal, may turn an initially positive experience into a source of anxiety or frustration, especially when users fail to meet the expected levels of engagement [Tuncay 2020, Loewen et al. 2019].

This article examines the design choices found within this platform's interface from a critical and ethical perspective to answer the following research question: *how may the design of Duolingo either promote user autonomy and clarity (bright patterns) or subtly manipulate child behavior (deceptive patterns)?* Our study reveals Duolingo's use of friendly pressure, such as the fear of breaking a streak, in addition to frequent and emotionally charged notifications and monetization tactics, which may obscure the children's agency. By analyzing how interface design reflects competing interests, such as commercial goals versus user well-being, a key contribution of this research is describing how persuasive strategies manifest in real use cases. We believe this study contributes to the ongoing debate on ethical design and encourages reflection on the balance between persuasion and manipulation in user experience.

2. Conceptual Background

This section introduces the two design paradigms that will be examined on the platform: the harmful approach of deceptive and manipulative patterns (Section 2.1), and the ethical and recent approach of bright patterns (Section 2.2).

2.1. Deceptive Design Patterns

Mobile applications have become an essential part of everyday life, fulfilling functions that span communication, education, entertainment, and commerce. However, their widespread adoption has brought with it a series of design practices that raise significant ethical and legal concerns, particularly when targeted at children [Valença et al. 2024]. Among these are deceptive and manipulative design patterns, which subtly steer user

behavior through the structure and aesthetics of app interfaces, often without the user's conscious awareness [Helamo 2023]. This issue is especially prevalent in mobile environments due to smaller screen sizes, persistent connectivity, and personalized notifications, which create conditions ripe for psychological manipulation.

In mobile apps, manipulation often takes the form of forced navigation paths, disguised advertisements, emotionally charged prompts, and fabricated constraints such as false urgency [Brignull 2024]. These features are not merely functional choices; they are calculated strategies designed to increase engagement, trigger impulsive responses, and ultimately increase profits, often through in-app purchases, recurring subscriptions, or the covert collection of personal data. These are practices that benefit companies financially, but compromise user autonomy and informed decision-making [Brenncke 2024].

Children are particularly vulnerable to these strategies due to their developmental stage and limited ability to recognize or resist manipulative intent. In mobile games, educational apps, and child-oriented social platforms, common tactics include offering in-game rewards in exchange for ad views, simulating scarcity to induce anxiety, or employing bright colors and sound effects to nudge children toward unintentional purchases[1]. These tactics, termed deceptive patterns, undermine conscious choice and encourage compulsive use, although often they are presented as features that "enhance user experience".

This concept describes interface designs that mislead users or trick them into actions they would not have taken otherwise. Brignull's taxonomy includes tactics such as preselection, obstruction, trick language, and visual interference, all of which are common in mobile applications [Brignull 2024]. In mobile-specific implementations, these may appear as default-checked boxes, ads disguised as game-play, or auto-renewal mechanisms buried deep within layered settings menus.

One notable example is *Confirmshaming*, where emotionally manipulative commands are used to pressure users into compliance. Questions like "are you sure you want to miss this opportunity?" or "don't you want to become smarter today?" appeal to guilt or fear, bypassing reasoned judgment. Similar techniques include *Fake Urgency* and *Fake Scarcity*, which employ countdown timers or warnings about "limited stock" to compel users to act hastily, without fully considering the consequences. In mobile contexts, such deception is intensified by limited display space, reliance on touch interaction and casual use, bypassing scrutiny and nudging users in moments of distraction or fatigue toward actions that do not align with their best interests. Hence, deceptive design obscures the true nature of platforms, such as excessive data collection and financial commitments.

Given this landscape, it is vital to expand the discourse around unethical design by revealing the unique features and risks associated with digital platforms. This is especially true for products aimed at children, where manipulative strategies compromise not only their experience but also core rights. Identifying, classifying, and ultimately regulating these design patterns is a necessary step toward creating digital ecosystems that respect user agency and promote safe and ethical engagement, particularly for younger users

[1]Dark Pattern Games - https://www.darkpattern.games/

navigating mobile technology.

2.2. Bright and Fair Patterns

The idea of bright patterns is closely related to fair patterns, which focus on solving problems and empowering users and restoring user autonomy, and facilitating informed and voluntary decision-making [Potel-Saville e Da Rocha 2023]. Bright patterns refer to persuasive user interface elements that prioritize the user's well-being and long-term goals over immediate desires and business interests [Sandhaus 2023]. This concept, which emerged in response to the increasing criticism of harmful design practices, aims to address more subtle ethical dilemmas in technology design (particularly in automation, personalization, and recommendation systems). Rather than simply avoiding unethical tactics, bright patterns offer a positive framework that promotes sustainable user benefits. They are intentionally designed to resist the temptation of short-term gains when those gains come at the cost of user satisfaction. Instead of manipulation, bright patterns employ ethical persuasion to align user choices with their authentic interests and values.

These patterns are based on transparency and explainability. According to [Sandhaus 2023], *Clear Consent* occurs when users are explicitly informed and asked for permission before their personal data is collected, with easy-to-understand language and visible opt-in or opt-out choices. In its turn, *Honest Defaults* also stands out by offering default options that reflect the user's best interest, such as not pre-selecting marketing subscriptions. *Informational Summaries* similar to nutritional labels provide users with a clear overview of the potential consequences of an action, such as how artificial intelligence features use specific data. *Persona Profiles* informs users about how they are categorized based on their behavior or data. Transparent recommendation systems disclose the logic or criteria behind content suggestions, building trust and clarity.

Another key element is making complex technical processes understandable. *Simplified Explanations* help users understand how their data is processed, while *Data Trails* shows what information is being stored, accessed, or shared, enhancing accountability and *Cost Transparency* breaks down the actual costs behind a product or service, fostering fairness in digital marketplaces. Bright patterns can also counteract digital echo chambers. *Out of My Bubble* exposes users to alternative perspectives on specific features, challenging their usual views. Similarly, *Contrarian Companion* introduces critical or contrasting opinions related to the content with which a user interacts.

Bright patterns may also be promoted via policy initiatives. This requires collaboration between lawmakers and design professionals to set clear definitions and create legal frameworks that support responsible design practices. This is the best strategy to establish a proactive and ethical approach to digital solutions, placing users (mainly, children) welfare at the center through transparency, clear communication, and long-term thinking.

2.3. Children's Well-being and Autonomy

The digital environment offers numerous opportunities for children and adolescents, but it also presents risks that can lead to rights violations and negatively impact their full development. Ensuring their well-being online is directly linked to the promotion,

respect, protection, and fulfillment of their rights. The United Nations Committee on the Rights of the Child developed General Comment No. 25 to guide states on measures to ensure digital rights online, in line with the Convention on the Rights of the Child [Livingstone et al. 2024]. In Brazil, the Federal Constitution (Article 227) and the Statute of the Child and Adolescent (ECA) establish the absolute priority and comprehensive protection of the rights of children, a principle that also extends to the digital context.

Children and adolescents are particularly affected by deceptive patterns and are considered a vulnerable population online. One third of Internet users worldwide are minors who require greater protection under regulations such as the GDPR (General Data Protection Regulation) [Chamorro et al. 2024]. Teenagers, despite their specific position and understanding of the online world, may be more likely to take risks, but also lack the resources to cope with potential harm. Their limited digital literacy makes children and adolescents more susceptible to the persuasive effects of online advertising and less aware of the complexity involved in the collection, analysis, and use of data for commercial purposes [de Oliveira et al. 2024].

Furthermore, they are in a stage of psychomotor development, including emotional and cognitive skills, making them vulnerable to manipulation of their online activities, consumption habits, health status, preferences, and other aspects of their developing personalities through the use of their data [Núcleo de Informação e Coordenação do Ponto BR 2024]. Hence, the digital autonomy of children and adolescents becomes a fundamental aspect of their online experience [Valença et al. 2024]. A reflective design approach can promote individual well-being by acting as the conceptual and practical antidote to manipulative patterns. In this paper, we aim to foster debate on fair design for child users, including in learning environments.

3. Research Method

To contribute to the ongoing debate on ethical design in language-learning platforms, we conducted a structured analysis of Duolingo. Our investigation followed a multiphase approach, combining exploratory usage with theoretical comparison.

3.1. Phase 1

We initially conducted an **exploratory use of Duolingo** with the goal of collecting preliminary data to support the subsequent stages of our research. This approach was qualitative and observational in nature, aimed at identifying interface elements, interaction strategies, and potential design patterns with manipulative characteristics. We began with the hypothesis that Duolingo did *not* employ manipulative design practices intentionally - i.e., we assumed the app sought to engage users ethically, without relying on techniques that exploit cognitive vulnerabilities to drive retention or conversion.

To verify this hypothesis, two researchers independently participated in this exploratory phase by using the Duolingo app continuously for a period of 10 days. We created new user accounts, ensuring an initial experience free from preexisting data or previous usage influences. In addition, each researcher selected a language with which they had no prior familiarity, in order to simulate a typical new-user scenario of discovery and learning. This setup aimed to reproduce a new user's journey, reducing experiential biases and enabling more neutral observation of the platform's interactions. The use of

two distinct participants also facilitated data triangulation, enhancing the reliability of the observations and minimizing the influence of isolated subjective interpretations.

We began this phase by engaging with the most common user flows, particularly those immediately presented on the home screen. These flows included core learning features such as listening, writing, speaking, and reading exercises, which form the pedagogical structure of Duolingo. We analyzed these activities in terms of their structure, visual and textual feedback, and the incentive mechanisms used to promote progress. Then, the focus was directed to the gamification features integrated into the platform. In particular, we examined functionalities such as the "leagues", which involve weekly rankings where users compete with others based on the number of experience points accumulated. This feature introduces a competitive dynamic into the learning process, which can significantly influence user behavior and engagement levels.

Moreover, we conducted a detailed assessment on the premium subscription plan offered by Duolingo by analyzing the advantages promised by the paid version, such as unlimited lives (allowing uninterrupted practice regardless of mistakes), the complete removal of advertisements, and the ability to freely revisit lessons. We evaluated these elements based on manipulative strategies that might encourage users to move from the free to the paid version. Finally, we explored the settings and user preferences available within the app. This included verifying options such as canceling subscriptions, disabling notifications, and accessing privacy and data control information. These aspects are relevant for understanding the level of autonomy users have in managing their experience and controlling external stimuli promoted by the app itself.

This simultaneous exploration of the platform ensured greater reliability in the observations and minimized subjective interpretation, allowing us to raise the conclusion validity of this research [Zhou et al. 2016].

3.2. Phase 2

In this phase, we performed a **literature review**, after completing the initial usage phase. This analysis was focused on *deceptive and manipulative patterns* and *ethical design principles*. Key references included academic studies and taxonomies on deceptive interface strategies, such as [Jarovsky 2022, Potel-Saville e Da Rocha 2023]. Based on this review, we defined a conceptual framework of design patterns against which we could evaluate features from Duolingo. The framework included categories such as *Obstruction*, *Nagging*, *Forced Action*, *Social Comparison* and *Fake Scarcity* [Brignull 2024]. From this stage onward, only one researcher continued the analysis.

3.3. Phase 3

During phase 3, we conducted a comprehensive **comparative analysis** by systematically cross-referencing the empirical data collected from the Duolingo app with the theoretical framework developed in the previous phase. The goal was to identify and classify potential occurrences of deceptive or manipulative design patterns, as well as to highlight examples of ethically sound design choices.

Each feature or interface element identified during the exploratory use was examined in detail to assess if it exhibited characteristics consistent with established deceptive pattern taxonomies. This evaluation process was guided by predefined pattern

categories, such as *Obstruction*, *Nagging*, *Forced Action* and *Fake Scarcity*. This comparison phase included multiple methodological strategies:

- Continued interaction with the platform, in which the researcher extended the use of Duolingo to observe behaviors and design prompts not captured during the initial usage period;
- Annotation of screenshots, where visual elements from the app were captured and labeled with their corresponding pattern category, when applicable. We used these annotated images as evidence to support the identification of patterns;
- Exclusion of features that showed no relation to the patterns under investigation, such as the basic instructional content that followed standard pedagogical structures and did not display any persuasive or coercive design elements.

3.4. Phase 4

Finally, phase 4 focused on **reporting our findings**. As the comparative analysis progressed, we began drafting this paper. To ensure clarity and coherence, AI-assisted writing and synthesis tools (such as NotebookLM from Google[2]) were used for style and grammar revision, as well as for labeling and associating deceptive pattern typologies that we had found in the literature.

3.5. Use of Artificial Intelligence

Throughout the development of this article, artificial intelligence (AI) tools were used as a form of support in several stages of the writing and translation processes, aiming to enhance textual quality and facilitate scientific communication in a foreign language. The use of AI occurred in a complementary manner, serving as an ally in the revision and structuring of the content presented. Initially, AI was employed to perform spelling and grammatical corrections, as well as to suggest stylistic improvements in selected portions of the text. Subsequently, the translation of the article into English was supported by AI-based machine models. This phase was of particular importance given that the authors are not native English speakers, and AI assistance provided greater confidence in terms of linguistic correctness, textual fluency, and fidelity to the original content. It is important to emphasize that all suggestions and translations generated by artificial intelligence were thoroughly reviewed by human authors, with technical knowledge and expertise in the subject matter. We used these tools as writing support, not as sources of content, interpretation or refinement.

In Figure 1, we present a schematic image with the previous steps of the study to clarify the structure of the investigation.

[2]NotebookLM - `https://notebooklm.google.com/`

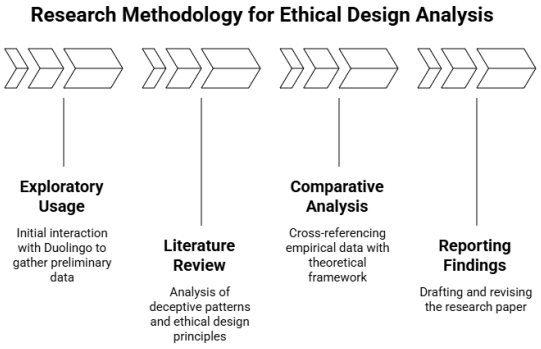

Figure 1. Four main steps performed by this research.

4. An Analysis of (Un)Ethical Design in Duolingo

In this section, we propose a critical analysis of the design patterns found on the Duolingo platform, with a specific focus on identifying and interpreting practices that may be classified as deceptive patterns or bright patterns. Our aim is not only to describe such design elements, but also to reflect on the psychological impacts they may generate, especially among younger users. Based on our observations, we seek to foster a well-founded debate about the boundaries between legitimate motivational strategies aimed at promoting engagement in language learning and practices that, due to their intensity or frequency, may prove harmful to users' emotional well-being.

Our analysis is structured in three main parts. First, we offer an overview of the user experience within the app, highlighting the main components of the **learning journey** (Section 4.2). Then, we examine the identified **patterns that can suggest manipulation or emotional interference** with the aim of increasing user retention (Section 4.3). Finally, we describe **ethical and constructive design practices**, which foster engagement in a transparent and positive way (Section 4.4). This structure provides a more balanced and contextualized understanding of Duolingo's interface and its interaction strategies.

4.1. Target Audience

Although Duolingo is frequently used by children, especially in educational and family settings, the app is not exclusively targeted at a child audience. Its core purpose is to serve anyone interested in learning new languages, regardless of age, educational background, or previous experience with digital platforms. The accessible interface, combined with gamified learning mechanics, makes Duolingo appealing to a wide range of users, from beginners to those seeking to improve existing language skills. In this context, children represent only one of the many age groups that make up the platform's user base.

The presence of colorful visuals, short challenges, and virtual rewards can encourage engagement among children, but these same features also help motivate teenagers and adults, making the learning experience more enjoyable and dynamic. Hence, Duolingo is designed for a broad and heterogeneous audience, including (but not limited to) children.

4.2. User Journey

After downloading and installing the Duolingo application via Play Store for Android and iOS devices, the user is automatically directed to the welcome screen. From this point on, the onboarding process unfolds in an intuitive and user-friendly manner, guided by a step-by-step configuration designed to personalize the learning experience. During this setup, we had to choose the language to study, indicating how we first heard about Duolingo, setting our daily learning goals (e.g., total of minutes to study per day), and deciding whether to add a Duolingo widget to our smartphone's home screen (a feature for quick access to progress tracking and reminders). Finally, the platform presents an overview of the expected outcomes after three months of consistent practice. Once the setup was completed, the application seamlessly transitioned us to the first learning activity, initiating the educational journey. These steps are shown in Figure 2.

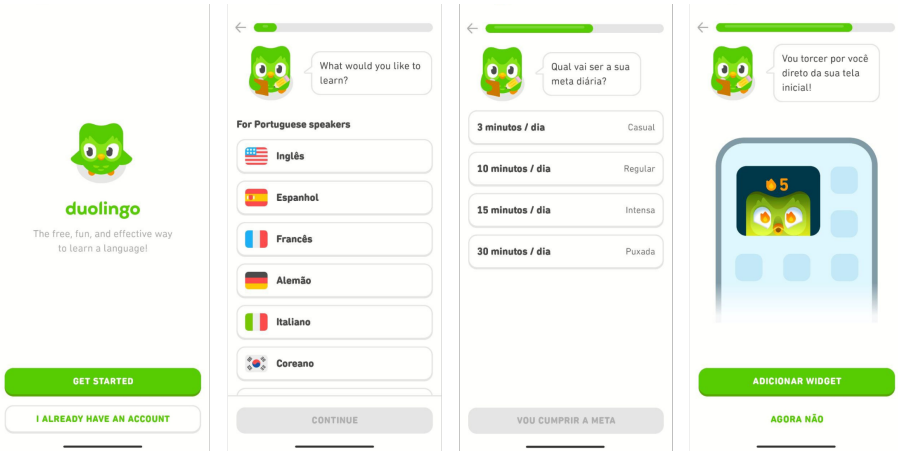

Figure 2. Initial interfaces that Duolingo presents to users.

Figure 2 shows 4 screens, where the user makes choices about their learning. In the (iii) image, the Duolingo mascot asks "What will be your daily practice?", while in the next screen he says "I will cheer you on right from your home screen".

The activities offered by the platform follow a structure based on small challenges in the language selected by the user, organized in a dynamic and interactive way. We noticed these challenges aim to (i) promote continuous learning and (ii) keep user engagement through a variety of exercises that simulate real-life language use. The main types of activities include: (1) *completing sentences* (an incomplete sentence is displayed in the target language with one or more blanks; users must select the correct word or expression from a set of options to complete the sentence); (2) *translating sentences* (a full sentence is presented in the foreign language, and users must translate it into our native language, demonstrating both semantic and grammatical understanding); (3) *listening comprehension* (a word or sentence is pronounced through an audio clip, and users must translate it accurately into our native language, exercising listening and oral interpretation skills); and (4) *pronunciation* (users must pronounce a sentence in the target language, when the app evaluates pronunciation based on phonetic criteria to encourage oral practice).

The platform uses a reward system to positively reinforce user performance. Each sequence of consecutive correct answers results in a congratulatory message and

contributes to a higher amount of *experience points* (XP) received at the end of the activity. The time spent completing the lesson also influences the final score. The faster the activity is completed accurately, the greater the XP awarded. Figure 3 shows an activity being carried out. In the first two screens, we can see the results of getting a lesson right and wrong, respectively. In the third one, we can read the text *"Each mistake takes away 1 life. You need lives to finish the lessons"*, informing the player about the number of attempts to get it right. In the next image, it is possible to see the possible ways to recharge the number of lives if the player reaches the limit.

Figure 3. Interfaces for lessons on Duolingo.

On the other hand, when an error is made, the user loses one of their five "lives". Additionally, their correct answer streak is broken, which may affect the total XP earned. This penalty system imposes a limit: once the maximum number of mistakes is reached (i.e., when all five lives are lost), the lesson is automatically ended. To resume, the user must restore their lives, which can be done in two ways: (i) purchasing the premium plan, which offers unlimited lives, or (ii) using gems, the app's virtual currency, which can be earned through in-app rewards or purchased with real money. If we choose to return to the home screen without continuing the lesson, we can earn lives by watching advertisements.

Upon successfully completing an activity, meaning that all challenges were answered correctly before all lives were lost, the user is taken to a summary screen showing the total XP earned and the time spent on the activity. Next, other screens are displayed showing progress on daily and weekly missions, highlighting completed and ongoing tasks, as well as the rewards earned, usually in the form of gems. If the user is participating in a ranking system (such as Duolingo Leagues), another screen shows their updated position on the leader-board after the XP has been counted. Between these transition screens, ads are common for users of the free version. From the home screen, Duolingo offers intuitive navigation to five other main areas of the platform (cf. Figure 4), each with specific functionalities, which are:

- *Speaking and listening lessons*: on this screen, the user can exclusively focus on improving their listening comprehension and speaking skills.
- *League ranking*: Here, the user can track their accumulated score and current position in the league they are participating in.

- *Profile*: this section gathers various pieces of information and customization options. The user can edit their avatar, view data about their performance and language, learning progress, such as the number of consecutive study days, and also access the settings screen to adjust preferences related to their account, notifications, privacy, and other technical aspects.
- *Quests and rewards*: This area focuses on the goals provided by the platform, offering a clear overview of both ongoing and completed objectives. Upon completing them, the user receives rewards, such as chests containing gems.
- *News*: this screen presents additional content produced by the Duolingo team, such as short informative articles and fun facts related to languages and cultures.

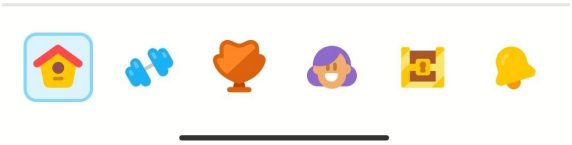

Figure 4. Navigation bar on Duolingo.

Finally, the Duolingo settings screen brings together various features that allow users to personalize their experience within the app. In this section, users can disable notifications, modify personal information such as their username and the email linked to their account, adjust language preferences, manage their subscription, and more.

4.3. Deceptive Patterns

Before proceeding with the following analysis of the deceptive patterns found in Duolingo, we must highlight that, depending on how the user responds to incentives, these patterns may not necessarily be deceptive or manipulative. A key part of this discussion is understanding what constitutes a beneficial incentive for learning and what, on the other hand, may be harmful to the user experience for child users. Even when the goal is language learning, we believe that these users are confronted with a deceptive pattern when design elements cross the line between motivation and manipulation, exploiting psychological vulnerabilities to force engagement or spending (such children's interaction with simulated gambling and in-game purchases, which cause addiction and can ruin those user's time and savings [Gupta et al. 2025]).

4.3.1. Cuteness and Interface Interference

A broader category of deceptive patterns known as *Interface Interference* involves manipulating the user interface in a way that favors certain actions over others, in order to confuse the user or limit the discovery of important possible actions. Within this category, *Aesthetic Manipulation* is one of the three identified subtypes, along with *Hidden Information* and *Pre-selection* [Gray et al. 2018]. Aesthetic manipulation can lead to misunderstandings about hierarchy, content type, or an unrealistic sense of urgency. It includes design choices that direct the user's attention toward one element in order to divert it from another. A specific example included in this category is "Playing with Emotion". This subtype encompasses the use of language, style, color,

or other similar elements to evoke emotion with the goal of persuading the user into a specific action. It may manifest itself through cute or frightening images, or appealing or intimidating language. Within this aspect, we analyze *Cuteness* design found in Duolingo, which employs visual and linguistic elements to tap into the user's emotions and influence their behavior [Gray et al. 2018, Lacey e Caudwell 2019], as shown in Figure 5.

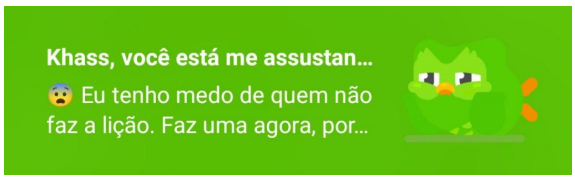

Figure 5. Example of notification exploring "Cuteness" deceptive design to invoke emotional responses and influence user's behavior.

This image reveals one of the most striking and strategically designed features of Duolingo: the use of visually captivating and intentionally "cute" characters. Its iconic mascot, the owl "Duo", is part of the app's overall aesthetic, which evokes a playful, friendly, and welcoming environment, conveying lightness and fun as part of its approach to making language learning more enjoyable and less intimidating. The characters (including Duo and a variety of human figures with diverse skin tones, hairstyles, clothing styles, genders, and ages) are consistently integrated into the user experience, appearing throughout exercises, notifications, and celebratory animations. Our analysis revealed that these excessive cute elements aim to foster emotional identification and affective bonds between users and avatars, as shown in Figure 5.

Figure 5 displays the message *"Khass, you are scaring me..."*. Just below it says *"I'm afraid you won't do your homework. Do one, now please."*. In another example, the character Duo changes based on how long the user takes to start a new lesson. This visual reminder encourages users to complete their daily lessons and may help keep children engaged and motivated in their language learning journey. However, the variation in Duo's expressions. In Figure 6, the icon shows the normal version of Duo in (a), a bored state in (b), and an angry state in (c). This can be seen not only as a persuasive technique, but as a subtle way of exploiting the emotions of young users. These children, who are still developing their self-esteem, can become undisciplined, late, or even irresponsible when they fail to meet the platform's expectations.

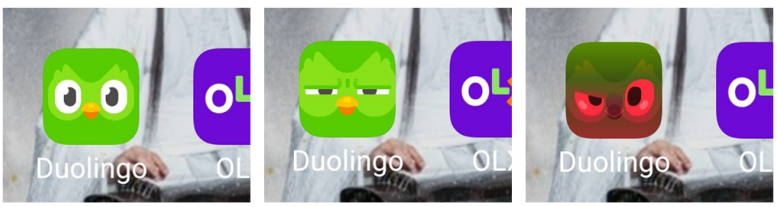

Figure 6. The icon of Duolingo (the character "Duo") exploring an emotional appeal in different stages in a single day.

This friendly aesthetic also serves as a vehicle for emotional pressure mechanisms. This becomes particularly evident in the "streak" system, which tracks the number of

consecutive days a user completes at least one learning activity within the app. When a streak is maintained, the user receives continuous positive reinforcement, such as encouraging messages, animated celebrations with the characters, and eye-catching visual indicators that show progress. On the other hand, when a user breaks the streak, that is, skips practice on a given day, the app responds with graphic and textual elements that may induce feelings of guilt or frustration. The mascot Duo, for instance, may appear with a sad expression, accompanied by phrases like *"We were doing so well..."* or *"Don't give up on me now!"*, suggesting a breach of commitment.

In addition, there are paid incentives to preserve the streak, such as the "streak freeze", which can be purchased using virtual coins or real money, allowing the user to maintain their streak even without practicing on a particular day. Although these strategies are effective in terms of user engagement, they raise questions about the extent to which the app's design respects user autonomy, or instead employs persuasive techniques that verge on manipulation, exploiting emotional vulnerabilities (especially among younger users) to ensure continued platform loyalty.

4.3.2. Obstruction

Another identified issue, which constitutes a more evident deceptive pattern, is known as obstruction. This pattern makes a process more difficult than necessary (which we can also call *user friction* [Mejtoft et al. 2019]) with the intention of discouraging certain actions [Potel-Saville e Da Rocha 2023]. It often appears as a barrier that prevents the user from completing a task they may want to do, but which is not a benefit for the platform (e.g. customize privacy settings for a more limited data collection). In the case of Duolingo, this pattern is primarily observed in how the app promotes its premium plan, when relevant information (especially information needed to make informed decisions) is intentionally made hard to access, requiring disproportionate effort from the user.

During regular use of the platform, the invitation to try the paid plan appears frequently, as mandatory ads lasting a few seconds and on specific screens that use bright colors and visual elements to highlight the "free" trial (cf. Figure 7). At first glance, this seems like a simple marketing strategy, but the lack of transparency regarding the actual subscription cost raises legitimate ethical concerns. In all instances where the paid plan is presented, the emphasis is on it being a free trial, often shown alongside a symbolic cost of R$ 0.00, without clearly informing the user of the real price after the free period ends. Even if the user chooses to close the ad or look for the information independently, such as navigating to the gem shop where the paid plan is located, the same pattern persists: the highlight remains in the "free trial" and the actual prices continue to be hidden.

Figure 7 shows an image with the title ''*Save with the Family Plan!"*. Just below it says *"Accelerate your learning in a group of 2 to 6 people"*. The button says *"TRY 1 WEEK FREE"*. To discover the real subscription cost, a user must go through six steps (cf. Figure 8): (i) first, they must select the free trial period; (ii) from there, they are directed to a screen outlining the benefits of the subscription; (iii) then, a screen shows a comparative table between the free and paid versions; (iv) after that, the user must confirm the trial once again; (v) then, the app informs the user will be notified two days before the free period ends; and (vi) finally, a screen shows the annual plans and respective prices.

Figure 7. Free plan invitation interface on Duolingo.

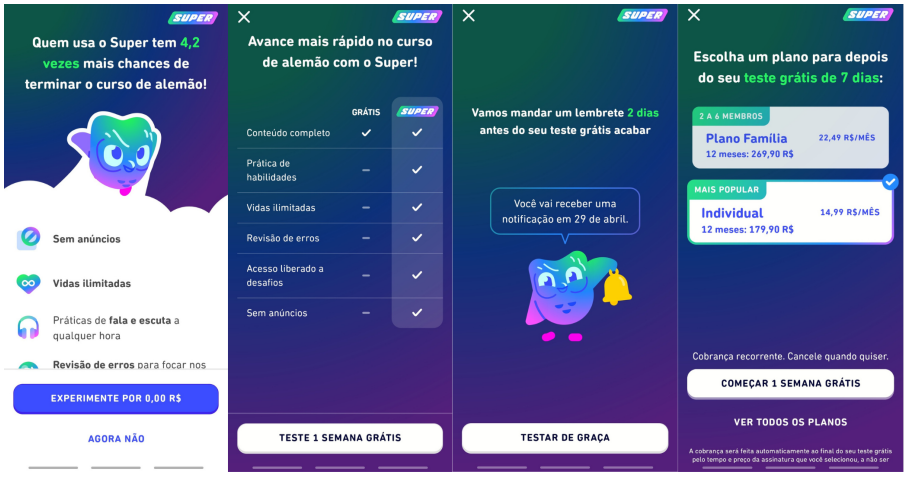

Figure 8. Example of *Obstruction* deceptive design.

Figure 8 shows 4 screens, where the user can read the benefits of the paid plan. The first screen has the title *"People who use Super are 4.2 times more likely to finish the German course!"*, then informs that the user will not need to watch ads, will have unlimited lives and will be able to practice speaking and listening at any time. The second screen continues to show other benefits of the plan. The third screen has the title *"We will send you a reminder 2 days before your free trial ends"*, the character then shows the message *"You will receive a notification on April 29th"*. On all screens, you can see buttons with texts such as *"TRY 1 WEEK FREE"* or *"TRY FOR FREE"*. Finally, the fourth screen has the title "Choose a plan for after your 7-day free trial:".

We observed that if the user wants to view the monthly plan, they still need to click *"see all plans"*, which leads to a new screen where the monthly option finally appears (highlighting that choosing the monthly plan will reduce the length of the free trial). Although the free trial is not automatically started before the final prices are displayed, meaning the user will not be charged without seeing the cost, the excessive number of steps and the intentional withholding of financial information until the very

last moment can create an environment of distrust. This exhausting process can confuse less experienced users, especially younger ones, and may lead them to make decisions without being fully aware of the financial consequences. Children and adolescents can, thereby, build an expectation based on the perceived benefits presented in the preceding screens.

4.4. Bright and Fair Patterns

Here, we describe some ethical design patterns in the user experience of Duolingo. As with deceptive patterns, the perception of these elements can vary depending on the user's profile and goals, what may be beneficial to some might not have the same effect on others. However, here we focus on aspects generally considered as positive design practices.

4.4.1. Escape Hatch

When looking at other apps, it is quite common to encounter significant difficulty when trying to delete one's account. In contrast, Duolingo stands out for offering a straightforward, transparent, and user-friendly process. The steps required to delete an account are minimal and clearly communicated. The user only needs to navigate through three screens before reaching the deletion option [Sandhaus 2023]. There are no misleading redirects or buried links, and the instructions are presented in a way that respects the user's time and decision. This approach can be classified as a bright pattern known as *Escape Hatch*.

This design pattern refers to features that make it easy for users to exit a service, cancel a subscription, or delete their profile. Rather than relying on obstacles to retain users, the platform provides an accessible route to departure, demonstrating a commitment to transparency and respect for user autonomy. In addition to offering a straightforward process for account deletion, Duolingo also allows users to cancel their premium subscription in a relatively simple manner. This feature is available in the app's settings area, where users can access the subscription management section. From there, they can view information related to their current plan and choose *"cancel"* if they wish to discontinue the service. Figure 9 shows these steps.

Although the flow involves a sequence of screens that could be debated from a user-centered design perspective, especially when considering standards that prioritize full transparency and clarity, the path to cancellation is neither hidden nor confusing. The button leading to subscription management is well positioned and the subsequent instructions are clear enough for users to understand what they are about to do, without visual traps or ambiguous language that could lead to mistakes (cf. Figure 10).

Figure 10 shows the screens for canceling the paid plan, called Super Duolingo. Screen (i) tells the user about the current plan and how long until the free trial ends: *"Free trial (6 days remaining)"* and *"Annual payments will begin on 05/14/2025"*. In blue, we have *"CANCEL SUBSCRIPTION"*. On screen (ii), the question appears: *"Why do you want to cancel Super?"*, showing options for reasons for cancellation just below. Among them are the options: *"Super's features are not worth the investment"* and also *"I signed up by accident"*, among others. Screen (iii) has the question: *"Are you leaving*

Figure 9. Interfaces for deleting accounts on Duolingo.

already?", followed by the text *"You will receive a reminder 2 days before your free trial ends!"*. Screen (iv) shows the message: *"If you cancel, the ads will appear again"*.

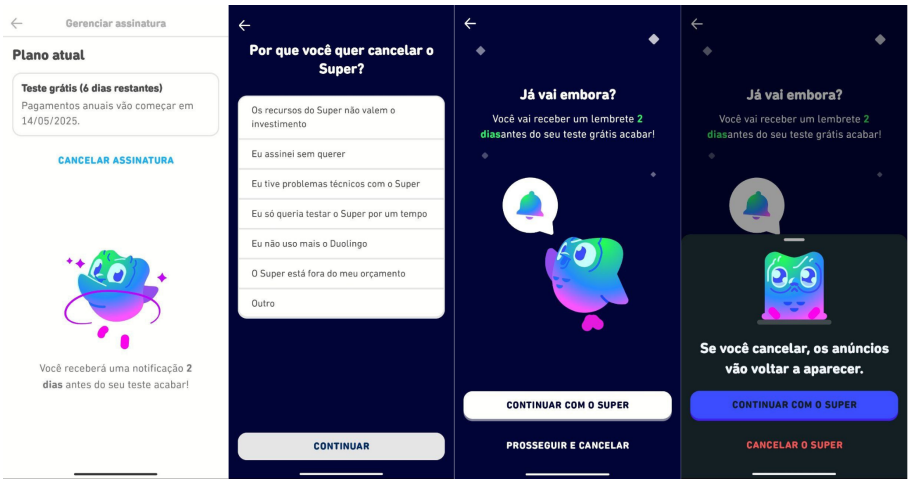

Figure 10. Interfaces to cancel a subscription on Duolingo.

In a scenario where a child becomes uncomfortable with certain aspects of the app (for example, the high frequency of notifications or a change in personal preferences) and decides to delete their account along with all associated data, Duolingo facilitates this action. Hence, given the no unnecessary friction, this feature fosters trust and serves as a positive model for other platforms. By empowering users to make informed and independent choices about their digital presence, Duolingo reinforces ethical design principles and shows that user satisfaction should take precedence over retention at any cost.

4.4.2. Fair UX

The pattern *Fair UX* refers to visual interfaces that respect user's intended actions and choices. These interfaces are designed to be consistent and clear, presenting options in a

straightforward and easy-to-understand manner. Examples of how Fair UX honors user intentions include the position, shape, size, and prominence of buttons, as well as the clarity of meaning in buttons and icons. The goal is to ensure that the visual interface aligns with what the user truly intends to do [Potel-Saville e Da Rocha 2023].

Duolingo's user experience stands out for its simplicity and effectiveness. Navigation is intuitive, with all the main features in a logical and accessible presentation, without visual traps or design decisions that mislead or confuse children. From actions like customizing an avatar to settings adjustments, several features require just a few clicks to be accessed. For example, changing preferences requires only one additional screen compared to avatar customization, maintaining a smooth and consistent flow.

A noteworthy aspect is that even features that are potentially unfavorable to the company, such as the option to disable all notifications, are clearly and easily available. This reflects a commitment to transparency and respect for user autonomy, prioritizing user experience over aggressive retention strategies. In addition, the language used in the interfaces also reinforces this positive perception: it is friendly and easy to understand.

Although this approach may seem basic, it is still rare among most applications. By offering a clean, predictable navigation without hidden obstacles, Duolingo fosters a greater sense of security and user control, strengthening trust in the platform. This clarity and respect for individual choice are key elements of human-centered ethical design.

5. Discussion

Critical analysis of Duolingo's design strategies reveals a recurring tension between motivational engagement and persuasive manipulation. Although the app incorporates practices considered ethical, such as intuitive navigation, positive gamification, and respect for user autonomy through patterns like *Escape Hatch* and *Fair UX*, it also presents concerning elements associated with deceptive patterns (cf. Table 1).

The repeated use of strategies such as the emotional appeal of the Duo mascot, the streak system, and the obstruction of pricing information suggests the user experience may be shaped by tactics that exploit emotional and cognitive vulnerabilities, particularly among younger users. This scenario highlights the need to balance commercial interests and digital well-being, especially on educational platforms aimed at vulnerable audiences. Rather than direct condemnation or unconditional approval, this research proposes a critical and thoughtful reflection on the role of ethics in user experience, emphasizing the importance of design standards that promote autonomy, transparency, and sustainable development of the relationship between users and digital platforms.

The ethical implications of these practices go beyond a mere annoyance. They directly challenge basic rights, especially when applied to children. The right to privacy, protection from commercial exploitation, and the right to autonomous decision-making in digital environments are all jeopardized by manipulative mobile design. This underscores the urgent need for regulatory frameworks that address not only deceptive content but also the structure and dynamics of the mobile user experience.

We believe that only by assessing deceptive and fair design patterns via evidence-based studies it is possible to develop regulatory solutions for ethical design to build digital platforms that ensure children's rights by design. Hence, we pave the

Table 1. Comparison of design patterns used in Duolingo

Observations	Escape Hatch	Fair UX	Cuteness	Obstruction
Description	Clear and accessible alternative to exit an action, without penalties or hidden obstacles.	Transparent design that respects user choices and promotes fair, deception-free navigation.	Use of cute or charming visual elements to soften frustration, create emotional attachment, and encourage actions that increase engagement.	Intentionally making certain processes more difficult or confusing to discourage actions like disabling notifications or turning off features.
Where it was found	Steps required to delete an account and cancel the premium subscription.	In the main features, with logical and accessible presentation, without visual traps or misleading design.	In the iconic mascot, the owl "Duo", and other characters presented during lessons or notifications.	Invitation to try the paid plan and the excessive number of steps to discover the real subscription cost.
Consequences	Makes the platform more reliable.	Greater sense of security and user control, strengthening trust in the platform.	May lead to emotional manipulation, compulsive usage, or attachment that overrides informed decision-making.	Can confuse younger users, and may lead them to make decisions without being fully aware of the consequences.

way to prioritize the rights and autonomy of young users, always respecting diverse cultural contexts (what is termed *multiple childhoods*) - something expected by the Brazilian Federal Constitution [de Oliveira et al. 2024] and reinforced by Plurality and Decoloniality in HCI [de Oliveira et al. 2024].

6. Final Considerations

In terms of **limitations**, we highlight that this study was conducted based on direct observation of Duolingo's interface and a theoretical review of the literature on design patterns. Hence, the research did not involved interviews, usability tests or controlled experiments with users, which could bring a relevant perspective of real experiences on deceptive design in particular. This approach may affect the generalization of our findings, being a threat to *external validity*. Furthermore, the analysis was based on a limited time sample (10 days of use), which may not account for future updates or variations in the app's interface, something that a longitudinal case would bring. However, the data

collected were accurate and precise, which improves the *internal validity* of this study.

In the theoretical field, a wider literature analysis (such as a mapping study), reviewing works in the area, would provide and involve the analysis of more studies addressing manipulative or persuasive design patterns in digital interfaces. This would include not only specific investigations on Duolingo or educational applications, but also broader research in the field of interaction design, user behavior, and ethics in technology. By assessing such complementary works, we would strengthen the theoretical framework of the research, allowing for a more critical and well-founded discussion of the findings.

Besides, interpretations of observed patterns are also subject to the authors' biases, as there was no triangulation with other empirical methods. To mitigate the risks of great inferences or inappropriate conclusions, we leveraged *conclusion validity* via multiple forms of evidence, notes taken while using the app and images that registered the identified deceptive or bright patterns, for example.

In **future studies**, we expect to conduct empirical studies with users of different age groups (e.g. separating children and adolescents, in order to better understand how they perceive and respond to the identified patterns considering different levels of knowledge and formation). In particular, we aim to adopt a participatory design approach [Bonacin et al. 2019] with children, obtaining their views on how to ensure ethical features - something reinforced by the dimension "Consultation" of the relevant framework *Child Rights by Design* [3]. In addition, longitudinal studies will assess the cumulative effects of exposure to manipulative or ethical elements over time. Finally, we aim to compare Duolingo with other educational platforms, mapping best practices and common critical issues. This study could contribute to the development of more robust ethical design guidelines for digital educational environments.

7. Ethical Care

This article did not involve empirical research with human subjects and did not include the collection of personal data or the execution of any actions that require interaction with participants. All analyses presented are based on theoretical review, observation of the platform's functioning in a public environment, and critical discussion grounded in scientific literature.

References

Bonacin, R., Dos Reis, J. C., e Baranauskas, M. C. C. (2019). Universal participatory design: achievements and challenges. *Journal on Interactive Systems*, 10(1).

Brenncke, M. (2024). Regulating dark patterns. *Notre Dame J. Int'l Comp. L.*, 14:39.

Brignull, H. (2024). *Deceptive patterns: Exposing the tricks tech companies use to control you*. Testimonium Ltd.

Chamorro, L. S., Lallemand, C., e Gray, C. M. (2024). " my mother told me these things are always fake"-understanding teenagers' experiences with manipulative designs. In *2024 ACM Designing Interactive Systems Conference, DIS 2024*, pages 1469–1482. ACM Press.

[3] 5Rights Foundation - *Child Rights by Design* - `https://childrightsbydesign.5rightsfoundation.com/principles/3-consultation/`.

de Oliveira, L. C., Amaral, M. A., Bim, S. A., Valença, G., Almeida, L. D. A., Salgado, L. C. d. C., Gasparini, I., e da Silva, C. B. R. (2024). Grandihc-br 2025-2035-gc3: Plurality and decoloniality in hci. In *Proceedings of the XXIII Brazilian Symposium on Human Factors in Computing Systems*, pages 1–19.

Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., e Toombs, A. L. (2018). The dark (patterns) side of ux design. In *CHI*. Pordue University.

Gupta, C., Campbell, M., Robards, B., e Fordyce, R. (2025). Playing the player: unfair digital gaming practices and their impact on australians. *Monash University*.

Helamo, A. (2023). Fighting the dark side: a scoping review of dark pattern mitigation. *University of Oulu repository*.

Jarovsky, L. (2022). Dark patterns in personal data collection: Definition, taxonomy and lawfulness. *Taxonomy and Lawfulness (March 1, 2022)*.

Lacey, C. e Caudwell, C. (2019). Cuteness as a 'dark pattern'in home robots. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 374–381. IEEE.

Livingstone, S., Third, A., e Lansdown, G. (2024). Children vs adults: Negotiating uncrc general comment no. 25 on childrens rights in the digital environment. In *Handbook of media and communication governance*, pages 414–428. Edward Elgar Publishing.

Loewen, S., Crowther, D., Isbell, D. R., Kim, K. M., Maloney, J., Miller, Z. F., e Rawal, H. (2019). Mobile-assisted language learning: A duolingo case study. *ReCALL*, 31(3):293–311.

Mejtoft, T., Hale, S., e Söderström, U. (2019). Design friction. In *Proceedings of the 31st European Conference on Cognitive Ergonomics*, pages 41–44.

Núcleo de Informação e Coordenação do Ponto BR (2024). Pesquisa sobre o uso da internet por crianças e adolescentes no brasil: Tic kids online brasil, ano 2024. Disponível em: `http://cetic.br/pt/arquivos/kidsonline/2024/criancas`. Acessado em: 25 de agosto de 2025.

Potel-Saville, M. e Da Rocha, M. (2023). From dark patterns to fair patterns? usable taxonomy to contribute solving the issue with countermeasures. In *Annual Privacy Forum*, pages 145–165. Springer.

Sandhaus, H. (2023). Promoting bright patterns. In *CHI*. Cornell University.

Shortt, M., Tilak, S., Kuznetcova, I., Martens, B., e Akinkuolie, B. (2023). Gamification in mobile-assisted language learning: A systematic review of duolingo literature from public release of 2012 to early 2020. *Computer Assisted Language Learning*, 36(3):517–554.

Tuncay, H. O. (2020). App attrition in computer-assisted language learning: Focus on duolingo. Master's thesis, McGill University (Canada).

Valença, G., Silva, J. V., Rocha, B., e Cortiz, D. (2024). Children's rights not deceptive patterns by design: a requirements perspective. In *Proceedings of the XXIII Brazilian Symposium on Human Factors in Computing Systems*, pages 1–11.

Zhou, X., Jin, Y., Zhang, H., Li, S., e Huang, X. (2016). A map of threats to validity of systematic literature reviews in software engineering. In *2016 23rd Asia-Pacific Software Engineering Conference (APSEC)*, pages 153–160. IEEE.