Politecnico di Torino

Computer Engineering

A.Y. 2021/2022

Graduation session December 2022

Impact of nudges on digital well-being for the smartphone

Supervisors:

Luigi De Russis

Alberto Monge Roffarello

Candidate:

Lorenzo Molteni

Abstract

In the recent past, a lot of research has demonstrated that an increasing amount of people, especially teenagers and young adults, is suffering from technology addiction. This new form of dependence can cause various health problems, starting from lack of sleep or decreased communication, and ending with depression, obesity and Attention Deficit Hyperactivity Disorder. To keep under control this problem, and possibly overcome it, different solutions have been experimented during the years. The most adopted proposal is represented by Digital Self-Control Tools, that are pieces of software, implemented as smartphone application or web extension, that help people in reducing their technology overuse by adopting interventions like timers, alarms, blocks of incoming notifications and screen-locks. Since DSCTs have been proven to not having stable and long-term effects, the newer concept of nudge has been introduced in the academic literacy. Nudges are alterations of the choice architecture that can affect the choice of users in a predictable way, without forbidding any of the options. Common examples in this context are choice defaults and larger buttons for the preferred option.

The main objective of this study is to investigate whether the nudge approach is effective or not, trying to give a response to the following research question: "Are nudges capable of driving the behavior of users while they are using social networks?". For reaching this goal, this thesis has been divided in three main steps. In the first one, six participants have been invited to a collaborative codesign session with the aim of designing the interface. Analyzing the results of this step, it has been possible to define both the list of social networks to be monitored by the interface, i.e., Facebook, Instagram and TikTok, and the types of intervention to be shown to user, i.e., a widget telling users whether they are scrolling fast or refreshing the page often, and a screen displayed before entering one of the targeted apps that informs about the mechanisms used in it to keep users online. In the final step, after the Android application has been developed following the design details previously specified, about twenty users have experimented the interface during a two-week long study.

By collecting the results of the user experiment, it has been possible to answer to the above-mentioned research question and, thus, to draw conclusions on the effectiveness of the nudging approach on the behavior of users while they are using social networks.

Acknowledgements

First of all, I would like to acknowledge my gratitude to Prof. Luigi De Russis and Alberto Monge Roffarello for their supervision and support during all the phases of this work. Their teachings, feedback, suggestions, and encouragements have helped me in accomplishing this thesis and reaching this important milestone of my academic career.

In second place, I wish to show my appreciation to all those colleagues and friends who have participated in the initial codesign session and in the final user study of this work, for the time they have devoted to this project without asking for anything in return.

A thought goes also to all those relatives and friends who are not here today and can not share with me this important day of my life. I am sure they would have been proud of me.

A special thanks is for Jessica, my partner, who was always right there when I needed someone to share my fears with. Having you by my side during this journey has helped me get back up after every fall.

Lastly, I give my deepest and more sincere thanks and gratitude to my wonderful family, my dad, my mom, my sister and my brother, that in the difficulties of life has found the strength to stay unite. Even if I am away from home, even if sometimes I seem further, you all are always in my heart and mind. If I have the chance to celebrate this goal today, part of the credit is surely yours.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The incredible technology developments of the last fifteen years, especially in terms of computational power inside mobile phones and network speed, have made it possible for software companies to design and develop more and more complex smartphone applications. Among all the others, a significant example is the one of social networks, which were firstly born as relatively simple website and are now available as really complex smartphone applications, having a huge variety of different features, from the possibility to send textual or vocal messages and share with others images and videos, to the capacity of providing real time, AI-powered photo filters.

The most famous social networks can be downloaded and used for free, thus this implies the existence of another way of making these projects economically sustainable and profitable for the company. These platforms attract the attention of their users, making them spend more time and engagement inside the application, and afterwards they sell this attention to advertisement campaigns for making their profit, in a business that has been named as "attention economy" [1]. By using mechanisms like dark patterns, which are further explained in the next chapter of this work, social networks are able to keep users online for a longer time, against their best interest and making them use the service even at times when they would not have used the service otherwise [2]. Unfortunately, on the other side of the display there are real humans, especially teenagers and young adults, and recent academic research has proven that, due to the above-mentioned mechanisms, the number of people suffering from technology addiction is constantly increasing. This new form of dependence can cause health disorders of different entity, from lack of sleep and decreased communication to depression, obesity, and ADHD.

Over the years, to overcome technology addiction and to reach the so-called digital wellbeing, intended as the correct balance that allows users to establish a healthy relationship with technology, some countermeasures were developed. One of the first and most used proposal was represented by Digital Self-Control Tools, which are, in short, pieces of software implemented as smartphone application or browser web extension with the aim of reducing technology overuse by applying interventions like screen-locks, alarms, passcodes and blocks of incoming notification. Examples of DSCTs can be found both inside the academic literacy [3, 4] and as commercial apps, e.g., Forest [5]. Despite this kind of tools have been proven to work well while the controlling system is up and running on the phone and interventions are shown to the user [6, 7], other studies have demonstrated that their efficacy in the long term tend to decrease [8]. Moreover, most of the current DSCTs gather external metrics, such as the overall time passed inside an application, being unable to detect, for example, passive and active usage of a social network, and they also strongly rely on self-regulation strategies that do not promote the habit-forming process.

In this scenario, academic researchers started exploring alternatives to traditional DSCTs, in order to overcome the limitations previously exposed. This thesis tries to apply the concept of nudge, which has been adopted in recent years by research studies in other fields, to the digital wellbeing area.

Nudges are defined as “any aspect of the choice architecture that alters people’s behavior in a predictable way without forbidding any options or significantly changing their economic incentives” [9]. According to their definition, nudges are clearly differentiated from DSCTs, since they are not bounded to technology, and they are also related to the concept of “libertarian paternalism.” This means giving users suggestions about the choice they should take, or at least the one that is considered better, but leaving them the freedom of choosing whatever they prefer. Thus, nudges are used to drive people in making better decisions and this approach can be also exploited by smartphone interface to fight against technology addiction. Some examples of nudges applied to the tech world are larger buttons for the preferred choice, choice defaults, and password meters that show the strength of the input by using a color scheme. Although nudges are a relatively new concept, many research studies have already demonstrated their efficacy, both in fighting technology addiction [10] and in other contexts, such as making better food choices [11] or opting for a pension plan [12].

1.1 Goal

This thesis has the aim of analyzing the efficacy of nudges in guiding the behavior of users toward more conscious decisions while they are using social media platforms on their smartphone. To reach this goal, it has been built a smartphone interface, called NudgeApp, for improving users’ digital wellbeing and, hopefully, diminishing their technology overuse. Its development has followed three main steps: the codesign of the interface, the definition of its details and contextual implementation as Android application, and, finally, a user experiment. Through the analysis of data collected during the above-mentioned user experiment, this work of thesis wanted to answer to the research question: "Are nudges capable of driving the behavior of users while they are using social networks?"

During the first step, six users have been invited to a codesign session. Here, through a user-centered and collaborative approach, they have shared their ideas on the list of smartphone applications to be monitored by the app and on the nudges to be implemented to discourage bad usages. After having analyzed the outcome of this session, the decision has been to monitor three social networks, i.e., Facebook, Instagram and TikTok, to focus on two specific dark patterns, i.e., infinite scrolling and pull to refresh, and to adopt two different nudges:

- **Nudge1**, which is a screen displayed before entering one of the targeted social media. It contains a short description of one of the dark patterns used in social media, together with possible diseases related to technology overuse, an animated gif that represents the DP in a graphic way, and two buttons, one for continuing to the app and the other to exit from it. This nudge has the goal of transferring knowledge to users and to bring them in a context where they are implicitly asked whether they are entering the platform for a real need or for just wasting their time.

- **Nudge2**, a circular widget containing an image that tries to catch the user attention whenever they are victims of one of the dark patterns of the social network. By clicking on the widget, it is possible to read why it has been shown. This nudge, instead, wants to make users understand when and why they are adopting a wrong behavior due to DPs and to trigger a reaction, e.g., start scrolling slowly and enjoy the contents, or stop refreshing compulsively the news feed.

In the second phase, the interface has been implemented as an Android application that runs a background service which listens to events raised by the operating system when the user interacts with the smartphone. The application, called NudgeApp, relies on an accessibility service that is capable of recognizing scrolls on the screen, detecting interactions that are considered wrong, and displaying the above-mentioned nudges. Bad usages are evaluated according to some metrics, e.g., number of pixels scrolled in the last thirty seconds of usage, time passed from the last scroll or scroll direction.

In the last step, around twenty participants have been contacted and invited to take part to a user study lasting two weeks. During the experiment, participants were first asked to download and install the application on their smartphone, and, then, to continue using their device as usual. Meanwhile, the background service started displaying nudges and sending information to a remote database, such as the timestamps of events like the display of a nudge and the opening or closing of a social, divided per application and per day.

Data collected through the experiment was crucial to evaluate the reaction of users to nudges and their behavioral changes during the study.

From the comparison of the two self-evaluation questionnaires about perceived smartphone addiction performed by all participants at the beginning and at the end of the test, it turned out that the developed interface helped more than one half of them in perceiving themselves as less addicted. Furthermore, by analyzing the usage data collected by the background service, NudgeApp showed to be effective in leading users towards a more normal and less compulsive behavior, thanks to the interventions proposed for the infinite scrolling and the pull to refresh dark patterns.

1.2 Structure of the thesis

The present work is organized as follows:

- The Background and related work chapter starts with an excursus on the notions that are the basis over which this work lies, where the reader is introduced to the concepts of dark patterns, digital self-control tool, digital well-being, and, finally, nudge. This background information is crucial to define the problem that this work is trying to solve, to expose the solution that has been developed in the recent past, to discuss over the new kind of intervention represented by the nudge, and, thus, to understand deeper the scope of this thesis.

- After this summary, I provide an overview of the codesign session in which six habitual smartphone users have designed the interface to be developed, defining a potential list of applications to be monitored and seven different possible interventions to be included. First, there is a preliminary description of what codesign is, including its transformation during the years and its theoretical basis, then, the goal, structure, and results of the codesign session are exposed.

- The next chapter, instead, embraces both the definition of the interface and the description of its development. Starting from the outcome of the codesign session described in the previous step, it is explained in detail which social networks are to be monitored, the type of interventions to be presented to the user, the metrics that govern when and how an intervention is displayed, and, finally, the interface architecture.

- The User study chapter is devoted at describing the experiment that has taken place in the last part of this work of thesis, which has involved around twenty social network users that have tested the developed interface for two weeks. During this user study, I have collected data relative to the interactions between the users and the interface and between the users and the monitored social networks. At the end of this chapter, the information collected from the user study are analyzed and the results are presented to the reader.

- In the last chapter of this work, I draw some conclusions on the effectiveness of NudgeApp for giving an answer to the research question of this thesis.

Chapter 2

Background and related work

Nowadays, almost all tech companies adopt a business model based on revenues coming from advertisements inside their websites or free-to-download applications. This implies that it is the best interest of the company to keep the user online as much time as possible, in order to maximize their income from advertising campaigns, in what has been named as "attention economy" by Davenport and Beck [1]. In this situation, the first question that comes to our mind is: how is it possible that companies, which don't even know who their customers are, are able to attract users' attention?

2.1 Digital well-being

The massive usage of electronic devices, from smartphones and personal computers to voice assistants and home automation, has for sure an impact on our everyday life. On the one hand, technology has opened an entire new world of interesting possibilities and benefits, and has helped in reducing social inequality by giving everyone the same and simple way for accessing crucial services, like education and healthcare [13, 14]. On the other hand, in recent years lots of research has pointed out the negative aspects of overusing technology [15, 16], and the possible implication of social medias in the growing rise of mental health issues, such as depression and anxiety, in adolescents [17, 18].

Over the years, technology became pervasive in everyone's life: from homes to workplaces, we are constantly surrounded by electronic devices, and our lifestyle, habits, and way of communicating and relating with ourselves and the others have changed accordingly, especially when digital technologies mediate our interaction with the surrounding environment. Just to give some significant numbers, users all over the world spend on the internet about 6 hours and 42 minutes every day [19], and about half of the teens in the US feel addicted to smartphones [20].

The adoption of digital technologies in critical social domains, like health and healthcare, education and employment, governance and social development, and media and entertainment [21], together with the related concerns and ethical issues, made the digital well-being topic gain popularity. On the web, and in the academic literature, there are plenty of different definitions for digital well-being [22], but, for sure, this term refers to the relationship between humans and technology and how this connection has to be manager for maintaining mental health and an overall sense of well-being. All in all, digital well-being can be seen as the correct balance that guide users in maintaining a healthy relationship with technology.

The increasing popularity of this topic favored the flourishing of third-party smartphone apps for limiting technology overuse. Recently, even tech giants like Google and Apple have published their own solutions to take this problem under control, which are, respectively, Digital Wellbeing [23] for the Android operating system, and ScreenTime [24] for iOS.

2.2 Dark patterns

The neologism dark pattern was coined in 2010 by the user experience designer Harry Brignull, who has also a doctoral in cognitive science, to address all those designs that were ethically dubious, meaning that they were not built that way by mistake, but they were “carefully crafted with a solid understanding of human psychology, and they do not have the user’s interests in mind" [26]. Since that moment, DPs have been studied by researchers all over the world, trying to build a common and exhaustive taxonomy for the most famous ones. The deceptive.design portal, formerly darkpatterns.org, established by the same Brignull [26], is a well-known example of taxonomy that shows samples of dark patterns on websites and mobile applications. Its famous “Hall of shame” is a collection where new examples of dark patterns are uploaded continuously by using the reports of many Twitter users, coming both from the real world and the digital one.

According to the definition above, and to the one provided by Gray et al., dark patterns can be seen as “user interfaces that intentionally manipulate people into performing actions against their best interest” [27]. This definition embraces a large variety of mechanisms, from trick questions and misdirection, to bait and switch, Figure 2.1, that are applied in almost all websites and applications at the expenses of the users.

An interesting subset of DPs is the one provided by Monge Roffarello and De Russis in their literature review [2], where the focus is on a particular kind of mechanisms called attention-capture dark patterns. This specific concept includes

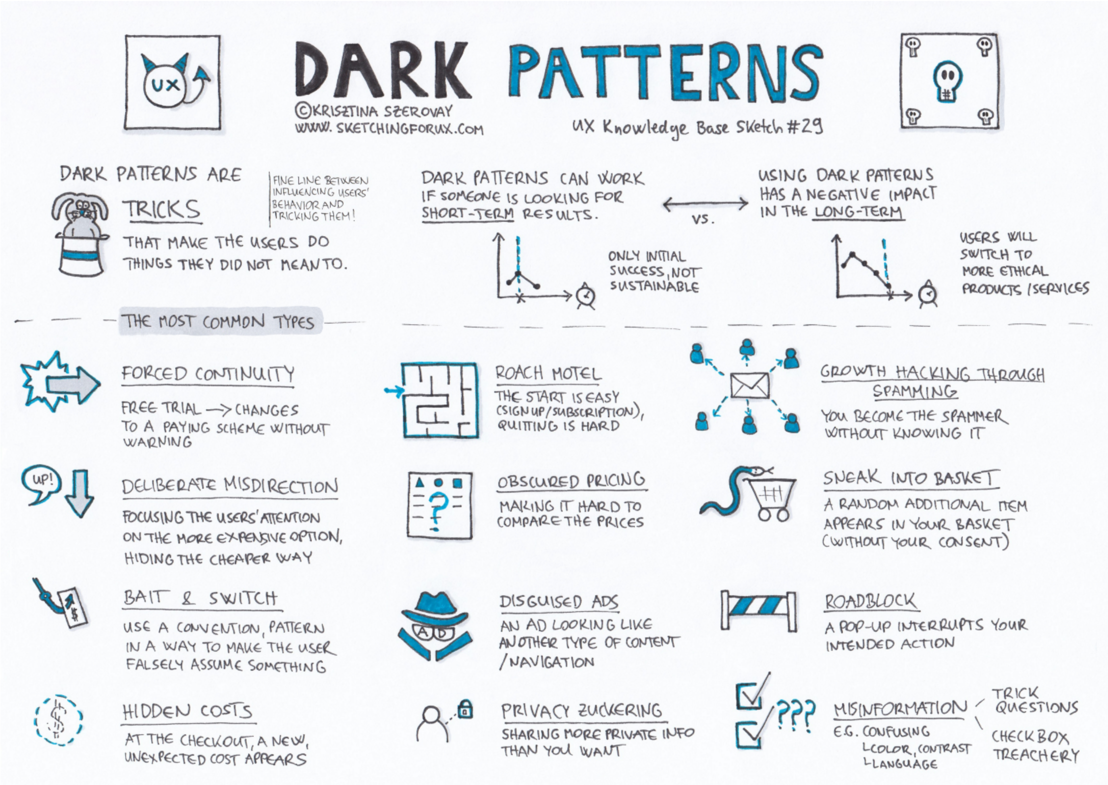

Figure 2.1: Dark Patterns
Source: Krisztina Szerovay [25]

all those malicious designs that are mostly related to technology, and social media, overuse, and that make people spend more time, attention, and engagement in a digital service, event at times when users would not have used the service otherwise [2]. In their work, after having reviewed the existing literature regarding DPs, and having performed an auto-ethnography, the authors have pointed out a list of 9 mechanisms, specific to the digital well-being context, that "consumes" the attention and time of users, e.g., notifications, recommendations, and gamification.

The concept of attention-capture dark patterns is particularly of interest for the goal of this thesis. Smartphones' applications, and in particular social networks, adopt these malicious designs for trying to maximize usage time, daily visits, and interactions/engagement, thus maximizing their profits from advertising companies, but this has implications on individuals' sense of agency [28] and in their lack of control over technology use [29]. Smart devices become a source of distractions, making users more stressed [30], and having, thus, a direct impact in their perceived well-being. The massive usage of these attention-capture mechanisms leads users become technology dependent and overuse it, adopting an incorrect behavior that can also cause problems like lack of sleep, decreased communication, depression, and also serious health disorders such as obesity and ADHD.

2.3 Digital self-control tools

Over the years, in order to overcome technology addiction and to reach the digital well-being, various proposals were presented. One of the most adopted solutions in this scenario is the one represented by Digital Self-Control Tools. DSCTs are a kind of software, which can be implemented both as a mobile application for smartphones or a web extension for browsers, whose aim is to reduce technology overuse by providing self-monitoring tools and countermeasures, such as screen-locks, alarms, timers, requests for closing a certain application, blocks of income notifications and so on. The interventions adopted by this kind of tools have been classified in four main categories [31], which are:

- **Block/removal**, features that allow to set limits, for example in terms of time spent inside a specific application, and provide blocking mechanisms, when the threshold is exceeded, e.g., locking the device or removing autoplay feature in streaming platforms.
- **Self-tracking**, features that show to the user his usage data, providing both an historical or aggregated visualization, and real-time feedbacks.
- **Goal advancement**, that make users set their objectives, in terms, for example, of maximum time to be spent on the smartphone per day, and then sends notification or comparison between the actual behavior and the desired one.
- **Reward/punishment**, that adopt an approach based on giving praise to users who performed well, and blaming the ones who had the worst behavior, for example by sharing the results among the community.

Digital Self-Control Tools can gather external statistics, such as the total time passed daily inside an application, but cannot collect, at least in an easy way, internal statistics about the specific different usages of each app. Thus, they cannot distinguish between the explicit intention of the user of using a social network for a specific and profitable purpose, and the “time consuming” usage. More specifically, websites and applications can be used for more or less legitimate uses: users can use a social media platform like Facebook both for getting in touch with friends or stay updated about a certain topic of their interest, and just for spending time scrolling down the newsfeed, falling into the “rabbit hole”[32] and getting stuck in those psychological mechanisms exploited by these applications to keep the user attention. Being unable to act from inside the application, for example by adapting the internal algorithm of a social network so that after a time threshold no new content is displayed, DSCTs can only implement higher level interventions, such as the ones described in the above list, and have no possibility of implementing more fine-grained restrictions, e.g., lock the screen only if the user is unintentionally using the device.

Despite DSCTs have been proven to work well as long as they are up and running in the monitored system, being able to intervene when needed, their efficacy in the long-term is yet underexplored [33, 34]. However, DSCTs are often not grounded in any underlying theory [35], and their interventions are not studied to trigger a behavioral change or a habit-forming process. Thus, as it has also been proved in a study of evaluation of 41 existing DSCTs [31], it is reasonable to think that their efficacy in the long-term decreases, after the monitoring tools is uninstalled or no more used, slowly returning to the levels present before the beginning of the interventions, unless other habit-forming strategies are used.

2.4 Nudges and boosts

In this scenario, it is clear that a new type of intervention is needed, in order to overcome the limitations imposed by traditional DSCTs. Thus, this thesis wants to explore the concept of nudge, which has already been introduced by academic researchers in other domains, such as health [11] and economy [12], and to apply it to the field of digital wellbeing.

A nudge is defined as “any aspect of the choice architecture that alters people’s behavior in a predictable way without forbidding any options or significantly changing their economic incentives” [9]. The concept of nudge is strictly related to “libertarian paternalism”, which means that the user is given the possibility of choosing whatever option he likes, since the choice is not constrained, but the design tries to influence him toward the right decision, for example making the desired choice the default, because, as stated by Sunstein, “nothing is what many people will do” [36].

According to the first definition above, and in order to clarify what a nudge is and what it is not, Calo has differentiated three distinct types of interventions: code, nudge, and notice [37]. A nudge is only what strictly matches the definition given by Thaler and Sunstein and reported in the previous paragraph, while a code is something that manipulates the environment having the goal of decreasing an undesired behavior, instead of increasing a desired one as done by nudges. Lastly, a notice is just information providing for users, as texts or reminders, which however previous studies have demonstrated not to reliably change the customer behavior [38]. Yet another type of intervention is that kind of influence called sludge, which

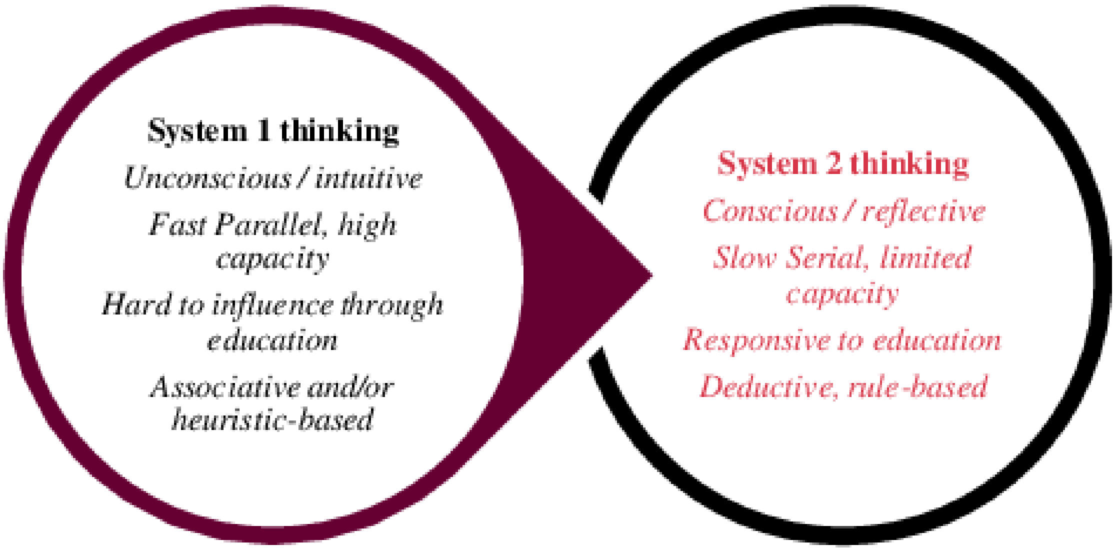

Figure 2.2: The Dual Process Theory
Source: Samuriwo and Pearman[41]

is, in short words, a nudge exploited by interface designers and companies to increase their profit margins conditioning the decisions of their customers [39]. An interesting study conducted in 2021 by Zimmermann and Renaud [40] has introduced a new kind of intervention, called hybrid nudge, i.e., the union between simple nudge and notice, proving that it performs better when compared to its single components and highlighting that increasing transparency, by providing additional information on how the nudge works and the reasons for which a certain behavior is encouraged, do not diminish the nudge efficacy.

On top of Dual Process Theories [42, 43], that are based on the underlying concept that humans have two distinct but connected cognitive processes, i.e., System 1 which is implicit and fast and System 2 which is rational and slow, Figure 2.2, nudges can be further subdivided. It is possible to distinguish Type 1 nudges, that primarily target System 1, such as choice defaults, and Type 2 nudges, that target System 2 after having activated in some way System 1 and trigger a reflective process, such as a password meter using color coding for displaying password strength that is able to attract the user attention via the color associated to the passphrase and, doing so, stimulates him to reason for increasing the complexity of the entered string [40].

Many academic studies have proven the efficacy of nudges [11, 10], even though their ability of triggering a behavioral change in the long term has not been well investigated yet. In this scenario, using hybrid nudges for an extended period, making the nudge itself less transparent but clearer by attaching an explanation, could help the habit-forming process.

In this same context, other researchers [44, 45] have studied the concept of boost, whose goal is to “train the functional processes or to adapt the external world, or both, to improve decision making and its outcomes” [44]. Differently from nudges, boosts aim at fostering users’ competence and knowledge about technology overuse and related problems, thus potentially having stable benefit effects in the long-term. However, it has to be said that most studies about nudge are short-term, so their efficacy in the long-term after having been user for an extended period has not been investigated. So, it could be possible that repeating a nudge over time causes a behavior change in users, according to what Hertwig and Grüne-Yanoff call “boosting side-effect” [44].

Chapter 3

Interface codesign

3.1 What is codesign?

“Co-design is about designing with, not for” [46].

The one above is a simple but significant definition that can help in understanding better the topic of this chapter. Codesign, also known as collaborative design, is a process where people’s goal is to build something from scratch, using a creative and participatory approach. This kind of collaboration has been exploited for almost forty years under the label “participatory design”, while the term codesign is much more recent.

This participatory approach is based on the idea that people having different backgrounds should collaborate within a design process [47]. The codesign process can be exploited in every design process, from simple to complex items, and requires both the participation of experts and non-experts. The first are people highly skilled in the specific field of the project, which are always crucial for taking into consideration every aspect that a non-expert can not know, such as functional requirements and limitation related to the domain, while the latter are usually represented by stakeholders, i.e., a member of “groups without whose support the organization would cease to exist” [48, 49]. Some typical examples of stakeholders, although this highly depends on the design process we are taking into consideration, are employees, system users, customers, and suppliers.

“In collaborative design, participants are not strictly bound to solve assigned partial problems but are encouraged to engage in solving design problems from other participants as well or to contribute to their design work” [50]. This other definition helps explaining the difference between collaborative and cooperative design: in fact, in the first case the team aims at solving the whole design problem as if it was a single person, while in the latter the team decomposes the design into smaller and simpler issues to be solved individually by someone in the group and then integrated. Since the whole codesign team has to act as a sole entity, it is thus clear that great communication is a focal point for the codesign process, together with some other key aspects such as diversity, support, sharing and motivation among team members [51], which are crucial to reach the desired outcome.

Collaboration during codesign can be classified according to two dimensions: openness and governance [52]. More specifically, in terms of openness a process can be:

- **Open**: everyone can join the team and contribute to the process exposing his own ideas.

- **Closed**: the participants, i.e., the members of the team, are chosen before the beginning of the codesign process by a manager or a group leader. With respect to an open organization, the team is less in number.

Whereas, in terms of governance the collaboration can be:

- **Flat**: participants make decisions together, after having reached an agreement, thus each of them has the same power inside the organization.

- **Hierarchical**: members are organized in a hierarchy, thus not all the participants have the same decision-making power and they have to deal with different challenges.

3.2 Goal of the study

This codesign study aims at building an interface for smartphones for increasing users' agency and consciousness while using social media platforms, thus improving their digital well-being and, hopefully, diminishing their technology overuse, through a user centered and collaborative design process.

The developed interface focuses on a list of social networks, or mobile applications in general, chosen by the participants of this preliminary study, and exploits the concept of nudge, introduced in Chapter 2, for improving the users' awareness of attention-capture dark patterns, their sense of agency, and their consciousness while using the smartphone. The introduction of the nudging concept inside the interface could happen, according to users' choices, in many ways. Just as an example, this can mean displaying a screen before entering the selected app showing the list of dark patterns present in the social media the user is going to interact with, or attracting the user attention when he is affected by dark patterns, specifying which psychological mechanism the app is exploiting and how. It is also important to underline that this interface has not the aim of changing the internal behavior of the selected apps, for example by disabling mechanisms like infinite scrolling, recommendations or autoplay, as it has already been done in other studies [53], but it wants to *add something to their interface*, by exploiting the nudge paradigm, trying to increase users' awareness of dark patterns and eventually letting them make more deliberate choices while using technology.

Thus, this session of codesign has a dual goal. On the one hand, it wants to understand which are the dark patters, among the presented ones, that users prefer to be informed of before entering or while using a social media. On the other hand, this collaborative study tries to discover the participants' preferred way of getting information about these mechanisms, for example by texts, images, videos or sounds.

3.3 Structure of the study

According to the classification taken from Pisano and Verganti [52], this codesign study has been carried out in a *"closed and flat"* fashion. This means that the group composition was predefined, with participants that were recuited from my personal contacts, and there was no decision maker, since everyone in the team had the same decisionsal power. More specifically, I played the role of group leader, being in charge of recruiting members needed for performing the experiment and coordinating the codesign session, without taking actively part in decisions. As previously said, the organization was flat, meaning that everyone in the group had to tackle the whole design problem and that decisions had been taken by reaching an agreement among all team members.

For the above reasons and for maintaining the group small enough to be easily coordinated and controlled during the design process, the number of participants in this preliminary study has been limited to 6. This study has been scheduled according to the users' need, both in terms of location and time. Thus, it has been conducted partially in presence with the first 4 participants, and partially online with the other 2. All participants were taken from my personal contacts and, in order to limit biases due to their backgrounds, the decision was to have half of them with previous technological studies and the other half with different educational paths.

Before the beginning of each of the two codesign sessions, all the participants were asked to fill in a **consensual form** and an **initial questionnaire**. Signing the first one, they have given their consent to collect and treat their personal information for the purposes of this thesis. The latter, instead, was useful to collect basic demographic information, i.e., age, gender, and occupation, and their general knowledge on the topics of this work, i.e., dark patterns, DSCTs, nudges, and technology in general.

This codesign process followed the four steps defined by the Design Council in

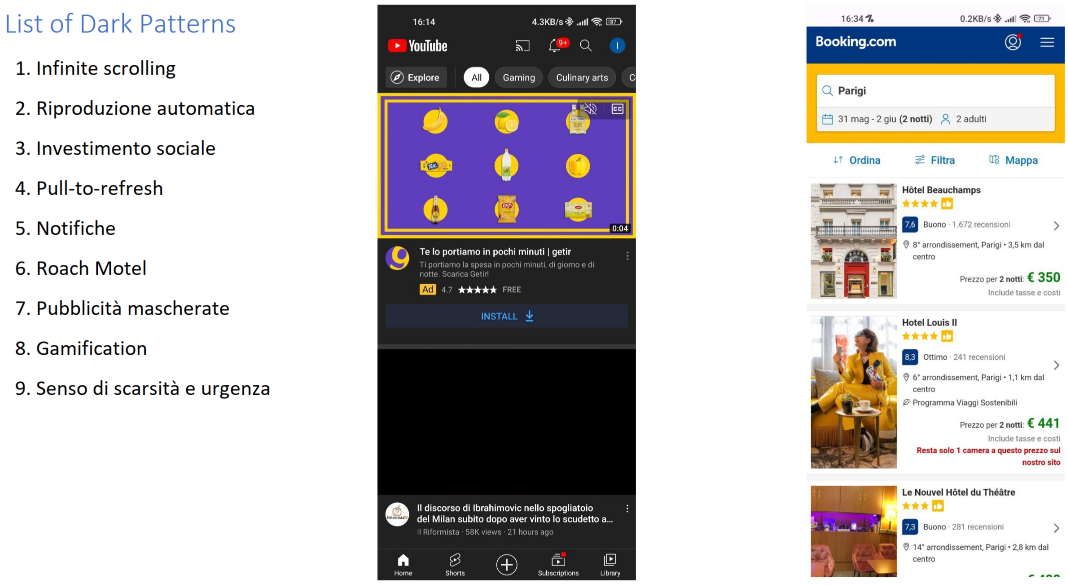

Figure 3.1: List of DPs - Examples of disguised ad and sense of urgency

their Double Diamond design model [54, 55], i.e., discovery, definition, development and deliver:

- In the **discovery** phase, which is aimed at scoping the work, participants were initially introduced to some basic concepts related to the field of this study, such as digital well-being, dark patterns, and nudges.
  First, the notions of digital well-being and nudges were presented joining their definitions with some simple real-life examples. Dark patterns, instead, were explained to participants providing them a predefined list of the most common ones, partially taken from the study conducted by Monge Roffarello and De Russis [2], and some images showing their usage in social networks (Figure 3.1). This list, and related images, has been available to participants throughout the study, helping them to keep clearly in mind what dark patterns are during the four phases.
  Secondly, users were introduced to the goal of this study, i.e., building an interface for increasing user awareness about dark patterns and improving their sense of agency. In this part, it has been clarified that the interface is intended for smartphones only, and that its goal is not to change the internal mechanisms of the selected apps, but rather to add some feature that can let users make more conscious decisions while using social network.
  After this explanatory phase, participants were invited to make questions and give feedback about their degree of understanding of the goal and field area

of the study. Then, they were prompted to write down on a post-it a list of four or five applications they would like the interface to be applied to or that they think are the ones with the highest number of dark patterns. This list is one of the outputs of this design process, since all the written proposals have been analyzed before the interface development to decide the applications to be monitored.
Finally, there was a dialogic conversation, aimed at diving deeper into the "world and vision of others" [47], where each team member could come up with possible issues and/or solutions regarding the interface to be developed. Here, participants had the possibility of proposing their own ideas about how to build the interface , for example on which strategy to use for attracting the user attention or on how to design the nudging mechanism for increasing their awareness. Furthermore, although the list of dark patterns was fixed, team members could still expose other kind of problems that affect, in their opinion, the usage of mobile applications and make them spend more time inside the application, against their best interests.

- In the **definition** phase, all the possibilities proposed in the previous step were analyzed and interpreted by the team, with the final goal of reaching, through a conversation-based approach, a convergence upon a specific result. This process has encountered an initial struggling, due to some divergences on the interventions to implement, but it helped also in identifying further problems about the proposals and hidden opportunities.

- In the **develop** phase, the team started creating the interface. This creative step followed, once again, a collaborative and cooperative approach, rather than a competitive one, where dialogue has been encouraged, and every opinion has been taken into consideration. Participants in presence had to design the low-fidelity prototype of the interface chosen in the definition phase using pencils and sheets of paper with the printing of some frames of mobile phone. For the ones online, instead, the low-fidelity prototype has been drawn by me, the group leader, following the descriptions about the interface given by participants.

- In the **deliver** phase, the produced paper prototype was analyzed and finalized, so that the outcome of this four steps process correctly matched the goal of the study. This implied the active participation of all team members to a final debriefing session, useful for getting feedback on the overall level of users' satisfaction regarding the final result.

It is however important to emphasize that this codesign session, as every other that is based on the Double Diamond design model, (Figure 3.2), was not a completely linear process that went straight through the four phases, since the

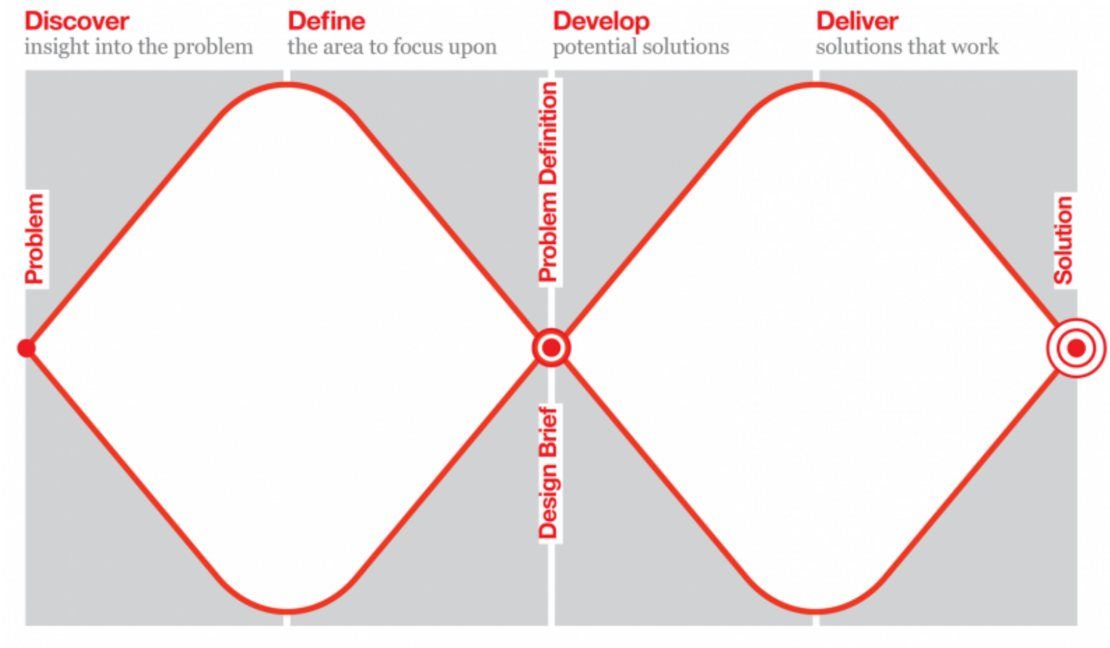

Figure 3.2: The Double Diamond model
Source: Justinmind[54]

results of each of the phases described above have been analyzed and taken into consideration to possibly rethink something coming from the previous steps.

3.4 Study results

This codesign study has been carried out in two separate sessions: the first one in presence with four participants and the second one online with other two participants. The six users, four males and two females, had a mean age of 26.67 years (standard deviation 10.98), three of them were master's degree students in computer engineering, one was graduated in statistical sciences, one was a law student and the last was a history student. Their answers to the initial questionnaire are reported in Table 3.1.

One participant out of six declared to be non expert in technology, but none of them had ever heard something about dark patterns or nudges. Concerning digital self-control tools, just one user stated he knew what they are, while the others didn't know. None of them was unsure about the answers to give.

When requested to list some of the applications they use more, in terms of daily accesses or total time spent, or the ones where they found the major number of dark patterns, the participants answered as follows in Table 3.2, where results have

	Really experienced	Experienced	Quite experienced	Inexperienced	Not expert
How expert are you in technology?	1	2	2	0	1
	Yes	Not sure	No		
Do you know what Dark Patterns are?	0	0	6		
Do you know what DSCT are?	1	0	5		
Do you know what Nudges are?	0	0	6		

Table 3.1: User responses to the initial questionnaire

been ordered by number of occurrences, and applications that were mentioned by only one participant were not considered. As it can be seen from the table, in the first three positions, two are occupied by social networks in the strict sense, i.e., Instagram and Facebook, while the other is the YouTube video platform. Interestingly, five participants out of six have mentioned Instagram.

	Instagram	YouTube	Facebook	TikTok	Netflix	Clash Royale	WhatsApp
Occurrences	5	4	3	2	2	2	2

Table 3.2: Users most mentioned applications

During both sessions, the list of dark patterns was predefined, but participants had been requested to expose any other problem that they had encountered during their smartphone usage, and they thought it was attention-capture. While in the first session none of them has exposed further problems, in the latter a user recognized that the presence of a player ranking inside a video game made him play more for reaching the highest positions and earn in-game prizes. However, ranking mechanisms can be included in the broader categorization of dark patterns related to social investment, which includes all those mechanisms that "influence users by instilling the idea that they should continue to use the platform to avoid loosing the achieved progresses" [2].

The following list contains all the proposals of intervention resulting from the two codesign sessions. Ideas from one to four are relative to the first session, while the others come from the second one.

1. Notification that arrives only if you are using your smartphone and informs the user, in a textual way, about dark patterns and/or a specific dark pattern. This notification is delivered once or twice a day, and the user has the possibility to enable/disable it.

2. Notification or popup message that, when the user has lost time due to a dark pattern, tells him that he is using much the smartphone and explains him the dark pattern that is affecting his behavior. In order to define what "much" means, the interface considers both the cumulative time passed from the user inside the application and the single session duration. The user can choose if

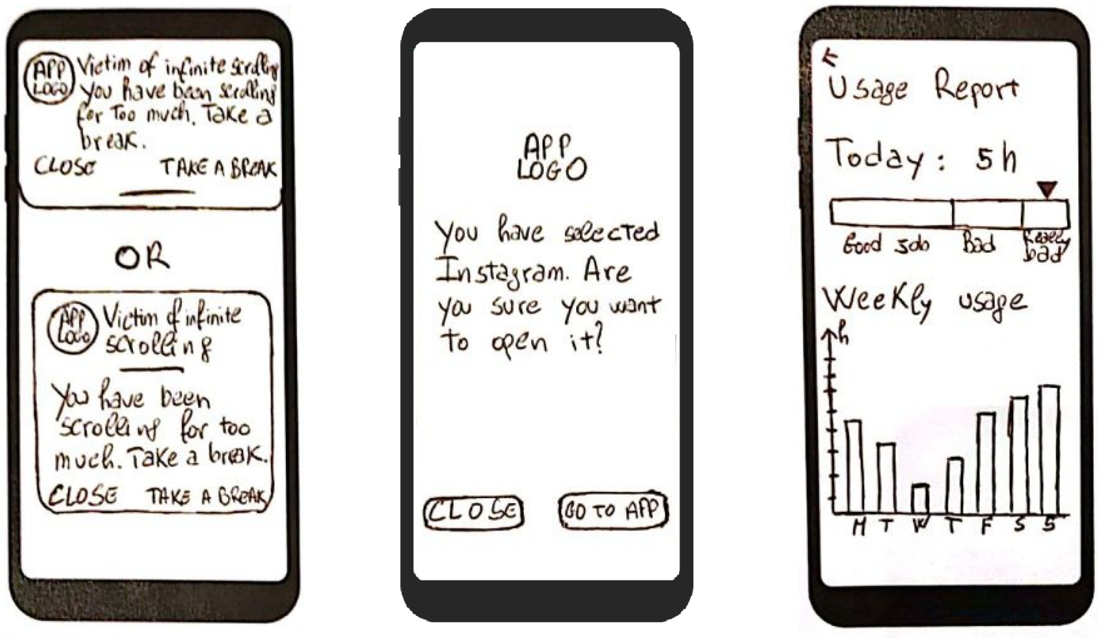

Figure 3.3: Phone frames - Nudges 2, 3 and 6

he prefers a notification or a popup message and can modify the text shown. This message is automatically shown again after some time if the user has not closed the app, but this function can be disabled.

3. Screen that is shown before entering an application, containing the number of daily accesses to the app, a predefined message or a user-defined one, and two buttons with customizable texts, one for entering the app and one for closing it. This feature can be enabled/disabled.

4. After the message telling the user he has lost time due to a dark pattern (point number 2 of this list) decrease the internet connection speed so that contents are loaded slowly. This function can be disabled.

5. Notification on screen shown after thirty minutes of continuous usage of an application telling the user to take a break.

6. Statistical report visible for the user showing daily usage and weekly average.

7. Notification on screen when the user has used the smartphone more than his weekly average telling him to take a break.

During the discussion about the proposal number 4, participants have also analyzed the possibility of decreasing the volume or the brightness of the smartphone.

However, after having discussed together, these two options were considered too intrusive, and thus discarded.

In the third phase of this codesign, called develop, after having established the list of nudges for the interface, users had to represent graphically, using a pencil and the provided phone frames, the interventions they had thought. Some of the ideas drawn, i.e., numbers 2, 3 and 6 of the above list, are depicted in Figure 3.3.

From the results of the four phases, i.e., discovery, definition, develop, and deliver, it is possible to make some interesting reflections regarding the dual goal of this codesign study.

In the first place, participants did not express a particular preference on the dark patterns they would like most to be informed about. This result could have been caused both by their general poor knowledge of the field area of dark patterns, clearly demonstrated in Table 3.1, and by their individual thoughts. During both sessions, i.e., online and in presence, all users focused on trying to build an intuitive and understandable interface, that could help in sharing knowledge about every specific dark pattern, or about the broader topic in general, without centering their attention in one of the mechanisms.

Secondly, by analyzing the list of proposed interventions, it is clear that the visual feedback is the users' preferred way of getting informed through the interface (6 out of 7 proposals). Specifically, four of them are notifications, sent by the interface periodically, when the user is caught by a dark pattern or when he is adopting a wrong behavior related to DPs, while the other two are entire screens showing more high-level information, e.g., a statistical report or other usage data. The only proposal not involving visual feedback is the decrease of the network connection speed. All in all, the study participants showed a clear preference in getting informed in a visual way, especially using asynchronous notifications or popup messages, while they almost never took into consideration other types of interventions, which were considered too intrusive.

At the end of this study, I made a preliminary analysis of the proposed ideas for identifying the ones that fit most the goal of this work of thesis. Since users clearly expressed their preference in getting informed in a visual way, the proposal at point 4 of the above list was discarded, i.e., the decrease of the network connection speed. Ideas number 5, 6, and 7, instead, were considered much related to traditional DSCTs, because they strongly relied on high-level information, like daily and weekly usage of smartphone, and they just provided types of interventions already implemented in self monitoring tools, i.e., notifications for telling users to take a break and statistical report.

Among the three remaining proposals, the first was asynchronous and not linked to a specific usage driven by those malicious designs. Thus, the choice was to focus on interventions at points 2 and 3, i.e., a notification shown when the user is victim of a dark pattern and a screen displayed before entering one of the time-consuming applications. These two nudges were the ones that required the direct intervention of users for being shown and that took most into account the influence of dark patterns on how they interact with their smartphone.

Chapter 4

Interface definition and development

This chapter describes the steps needed to define and develop NudgeApp.

After having analyzed the outcome of the codesign study, which have been discussed in the last section of chapter 3, this work continues with the definition of the interface details and its development. More precisely, sections 4.1 and 4.2 discusses the social networks monitored by NudgeApp and the two implemented interventions and the metrics that governs the display of the nudges, respectively, while section 4.3 focuses on showing some particular implementation details of the interface that can be of interest for the reader.

4.1 Interface details

After having carried out the codesign study with six habitual users, in two different sessions having four and two participants respectively, it is possible to analyze the results that directly affect the interface development. Looking at Table 3.2, containing the list of applications decided by the participants of the collaborative study, it's clear that most of them are social networks. As a matter of fact, among the first four most used applications, there are popular social networks like Instagram, Facebook and TikTok. Instead, the second most used app is YouTube, where users tend to spend much time mostly because of the length of the videos present in the platform.

For the above reason, this interface focuses exclusively on the social networks listed above, i.e., **Instagram**, **Facebook** and **TikTok**, which are the ones that participants used most or where they found the highest number of dark patterns.

Regarding the dark patterns to be analyzed by the interface, participants did not mention any preference on the DPs to be informed about. In this case, the

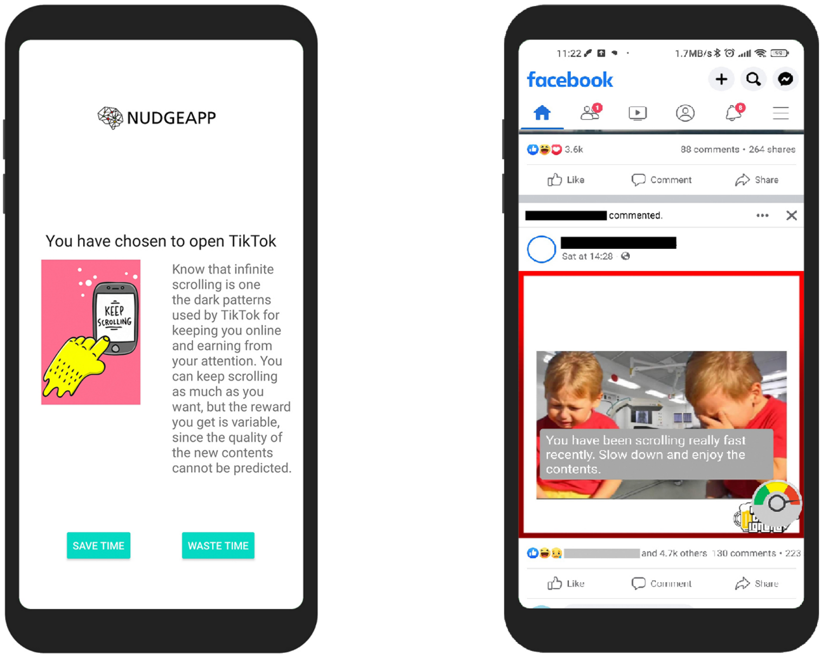

Figure 4.1: Nudge1 (left) and Nudge2 (right)

choice has been to develop an interface focusing exclusively on infinite scrolling and pull to refresh. These two attention-capture mechanisms were chosen because they require the direct interaction of the user, by swiping or tapping on the screen, while others like social investment, disguised ads and notifications exist inside most apps even if the user does nothing on the platform. Moreover, another reason that has driven the decision to focus specifically on these two dark patterns was that they were included in all the social networks monitored by this work.

Finally, even the decision on the interventions be implemented by NudgeApp was driven by the analysis of the results of the codesign study. Here, users showed their clear preference towards getting informed about nudges in a visual way (6 proposals out of 7). Then, among the six remaining ideas involving a visual feedback, interventions number 2 and 3 were selected and implemented in NudgeApp, while the others were discarded because they were already implemented in classical DSCT or because they just comprise asynchronous notifications. Moreover, the selected nudges were the ones that took most into account the direct influence of DPs on users behavior while they are using their smartphone and that were triggered by an explicit action. Before the development, both nudges were analyzed and further detailed in order to be directly implementable, as described in the list below.

1. **Nudge1**: A screen shown before entering one of the chosen social networks, containing a textual nudge, referred to one of the two dark patterns which

are the focus of this study, and two buttons, one for continuing inside the application and the other for closing it.

2. **Nudge2**: A notification, implemented as a widget image on screen, that appears in the bottom right corner whenever the user is inside one of the selected platforms and is victim of one of the two dark patterns the interface is focusing on, i.e., infinite scrolling and pull to refresh. The widget, which is a non-textual nudge, is accompanied with a textual nudge that is shown if the user clicks on the image.

Referring again to the Dual Process Theory [42, 43], both interventions are Type 2 nudges, since they primarily target the rational and slow System 2, after having attracted the user attention in some way, i.e., using the screen or the widget image.

4.2 Nudges and metrics

This section aims at diving deeper into the details of the implemented interventions, i.e., Nudge1 and Nudge2. First, there is a description of the textual and non-textual nudges used in both interventions, divided depending on their target dark pattern. After, they are explained the constraints and metrics, for example the scrolling speed for Nudge2 or the maximum number of daily appearances for Nudge1, that rule when and how both interventions are displayed.

4.2.1 Textual and non-textual nudges

Textual nudges are, essentially, short sentences that aim to transfer knowledge on DPs to users, helping them making more conscious decisions. In this work, two different types of textual nudges were developed for fitting better the contexts of Nudge1 and Nudge2. Particularly, for the first intervention, i.e., the screen shown before entering one of the monitored apps, textual nudges are longer. In this case, their goal is to inform the user about the dark patterns present in the social network he is going to open, the way in which they work, and the possible risks connected to wrong behaviors, like technology overuse, favored by DPs. Concerning the second intervention, i.e., the widget image that shows a textual nudge when clicked, the sentence is shorter, since it needs to fit a popup window without occupying the whole smartphone screen. In this case, the nudge aims at pointing out the wrong behavior adopted by the user, e.g., fast scrolling or compulsive page refresh, and at guiding him towards making better decisions.

Instead, in this work, non-textual nudges are used only in Nudge2 and are represented by images there are shown in a widget icon when the interface detects that the user has been victim of one of the monitored dark patterns. These images,

Infinite scrolling		
Textual nudge (Nudge1)	**Textual nudge (Nudge2)**	**Non-textual nudge (Nudge2)**
1. Infinite scrolling is one the dark patterns used by Facebook/ Instagram/ TikTok for keeping you online and earning from your attention. You can keep scrolling as much as you want, but the reward you get is variable, since the quality of the new contents cannot be predicted. 2. The infinite scrolling dark pattern used in social networks can lead you to a more passive and problematic technology overuse, causing lack of sleep, decreased communication, depression, or even obesity and ADHD. [56]	1. You have been scrolling fast/ really fast recently. Slow down and enjoy the contents. 2. You have been scrolling fast/ really fast recently. If nothing is of your interest, consider taking a break.	  

Table 4.1: Textual and non-textual nudges for infinite scrolling

which need to attract the user attention, are provided in three different versions representing three levels of interaction, from normal to compulsive, and they adopt the same color scheme of traffic lights, i.e., green, yellow, and red.

Table 4.1 describes in detail the nudges implemented for the infinite scrolling dark pattern. The first textual nudge for Nudge1 tries to explain the psychological mechanism over which this DP relies, while the second wants to highlight all the health problems that can be caused by the passive and problematic technology usage favored by these malicious design. Instead, both the textual nudges for Nudge2 focus the attention on the wrong behavior adopted by the user, i.e., fast or really fast scrolling, and provide a suggestion for improving the interaction between user and social network. As shown in the above mentioned table, for this dark pattern the choice was to use a speedometer as non-textual nudge, whose image has been retrieved from Dreamstime[1], having 3 different levels referring to normal usage, fast scrolling, and really fast scrolling.

[1]https://it.dreamstime.com/photos-images/tachimetro-3-livelli.html

Pull to refresh		
Textual nudge (Nudge1)	**Textual nudge (Nudge2)**	**Non-textual nudge (Nudge2)**
1. Pull to refresh is one of the attention-capture dark patterns used in Facebook/ Instagram/ TikTok. It allows you to swipe down and load contents which can be new or not, as it happens in slot machines, where you can win or lose. 2. The pull to refresh dark pattern used in social networks can lead you to use more technology, causing lack of sleep, decreased communication, depression, or even obesity and ADHD. [56]	1. You have refreshed the contents. If you have already seen everything, just take a break. 2. Pull to refresh is like a slot machine. This time it's a win or a loss?	

Table 4.2: Textual and non-textual nudges for pull to refresh

Table 4.2, instead, refers to the pull to refresh DP, and includes the details of the implemented nudges. Concerning textual nudges, both for Nudge1 and Nudge2, they were studied and written making the same reasoning explained before for infinite scrolling. In this case, the non-textual nudge is represented by a refresh icon, retrieved from Wikipedia[2], which uses the same color schema of the previous dark pattern, i.e., the green icon for normal behavior, the yellow one for frequent, and the red one in case of really frequent, or compulsive, usage of the feature.

4.2.2 Metrics

This section aims at describing the constraints and measures that have been thought for controlling when and how one of the interventions has to be displayed to the user.

Regarding Nudge1, i.e., the screen shown before entering Facebook, Instagram, or TikTok, it is important that the intervention is not displayed every time he is entering one of the targeted apps. In this way, the nudge itself avoids becoming

[2] https://en.wikipedia.org/wiki/File:Refresh_icon.svg

repetitive, intrusive, and annoying, and these three characteristics could also help users in feeling more in control. This is something participants really valued during the codesign phase, where they have chosen to have the possibility of activating and deactivating the feature in most of the proposals. Since all the applications taken into consideration are social networks, that are primarily used for interacting with contents like images, posts, and short videos, the defined constraints were applied in the same manner for all of them.

The one following is the list of metrics that has been implemented in NudgeApp for the first intervention.

1. Nudge1 appears at most once a day per each application, for a maximum of twice appearances per day. For example, it is no more shown if it has already been displayed once when entering on Facebook and once on TikTok. Since this intervention informs users about dark patterns, the limit on its daily appearances has been defined thanks to what participants said about their first proposal. In that case, they wished to be informed at most once or twice a day by a notification that could appear only if they were using their phone.

2. Nudge1 is displayed for application X, only if the user has already spent more than 30 minutes inside the social network X that day. The chosen amount, i.e., 30 minutes, is the average time that United States users spend daily on Instagram (29), Facebook (28) and TikTok (31), according to Insider Intelligence [57]. It is thus reasonable to consider showing the nudge only after the first 30 minutes of daily usage. Moreover, this is the same amount of time that the study participants have mentioned in intervention number 5.

3. Nudge1 is not displayed if it has already been shown for another social network within the previous 60 minutes. This avoids that the intervention is displayed twice in a short time interval, or even in two subsequent opening of two different socials, thus helping to spread better the appearances during the day and to keep the interface less intrusive for users.

4. When entering one of the targeted apps, the probability that Nudge1 is shown is 0.25.

5. If the user has spent more time than his daily average on the social networks in consideration, the probability that Nudge1 is shown is increased to 0.5.

For Nudge2, i.e., a widget image that appears whenever the user is victim of one of the two targeted dark patterns, the metrics have to vary depending on which DP is considered, i.e., infinite scrolling or pull to refresh.

Concerning infinite scrolling, as shown in Table 4.1, three different scrolling speeds were defined, i.e., normal, fast, and really fast scrolling. Thus, the metrics presented in the list below have to reflect these three possible behaviors, defining plausible unit of measures that can be monitored by NudgeApp.

1. Nudge2 is shown only after the user has scrolled the screen at least 10 times in that usage session. A usage session is defined as a period that starts when a social network is opened and ends when the application is closed.
2. The image shown inside the widget for Nudge2 varies depending on the number of scrolls detected by the interface during the last 30 seconds of usage. The time window is set to half a minute for allowing the user to recover faster from a situation of fast/really fast scrolling. Thus, Nudge2 is displayed as:
   - Normal scrolling, i.e., green speedometer, if less than 5 scrolls during the time window.
   - Fast scrolling, i.e., yellow, if the number of scrolls is between 5 and 9 (included).
   - Really fast scrolling, i.e., red, if 10 or more scrolls are detected in the last 30 seconds of usage.

   The above numbers have been calibrated according to my personal usage of smartphone, where I have emulated the three situations, from the first, in which I stopped and read the displayed contents, to the last, in which I just looked at pictures for a few seconds each.
3. After Nudge2 is shown, the image in the widget is updated according to the user interaction, passing from really fast to fast, or vice versa, and it stays on screen until the user's behavior does not return to "normal scrolling", according to the parameters defined above.
4. If Nudge2 is displayed on screen as "green" and the user maintains a behavior that is considered normal scrolling, i.e., less that 5 scrolls in the last time window, the widget disappears after 3 seconds. This measure helps Nudge2 in being less intrusive and annoying for the normal interaction between the user and the system.

The same high-level reasoning can be applied for the other dark pattern, i.e., pull to refresh. In this case, the defined metrics listed below take into consideration the number of gestures detected by NudgeApp during the usage session.

1. Nudge2 is not shown if the application has been opened less than 10 seconds ago. This constraint has been included because, as also confirmed by the study participants during the codesign session, users usually pull the screen when they have just opened a social for being sure that the page has been actually refreshed.

2. Considering as time window the last 10 minutes of usage, the image shown inside Nudge2 varies depending on the number of pull gestures detected by the interface. Thus, Nudge2 is displayed as normal (green refresh icon) if the user has pulled the screen once, frequent (yellow) if he has pulled twice, and really frequent (red) in case of three times or more.

3. After Nudge2 is shown, regardless of the picture it is displaying, it stays on screen for 3 seconds before disappearing, giving the possibility to the user to click on it and show the textual nudge.

Furthermore, another metric has been implemented for Nudge2, independently from the dark pattern it is referring to, i.e., infinite scrolling or pull to refresh. In particular, since the widget image can be clicked by the user to show a short textual nudge, in case this is displayed on screen the timer for hiding the widget is increased by two more seconds, bringing the total to five, for giving the user the possibility of reading the message.

4.3 Interface architecture

Before exposing the details of the interface architecture, it is necessary to explain some basic concepts related to the development of an Android application, which are essential to understand better the following of this section.

An Android application is made up of one or more components that have to be declared in its manifest file, which is an XML that includes the core information of the project and is used by the operating system to interact with the app. There are four types of components:

1. **Activity**, which is used to model a “single, focused thing that the user can do”[3]. Activities can create visual windows, i.e., the GUI where the user can interact with the smartphone, and can perform a meaningful task for the application. Android stores in an activity stack all the tasks that are currently running on the phone, and when one of them reaches the top of the stack it is displayed to the userm who can start interacting with it. Thus, only one task at a time can be foreground, represented by the activity which is currently on top of stack, while there is the possibility of having multiple background tasks.

2. **Service**[4], a component that runs in background and is ordinarily used to

[3]https://developer.android.com/reference/android/app/Activity

[4]https://developer.android.com/reference/android/app/Service

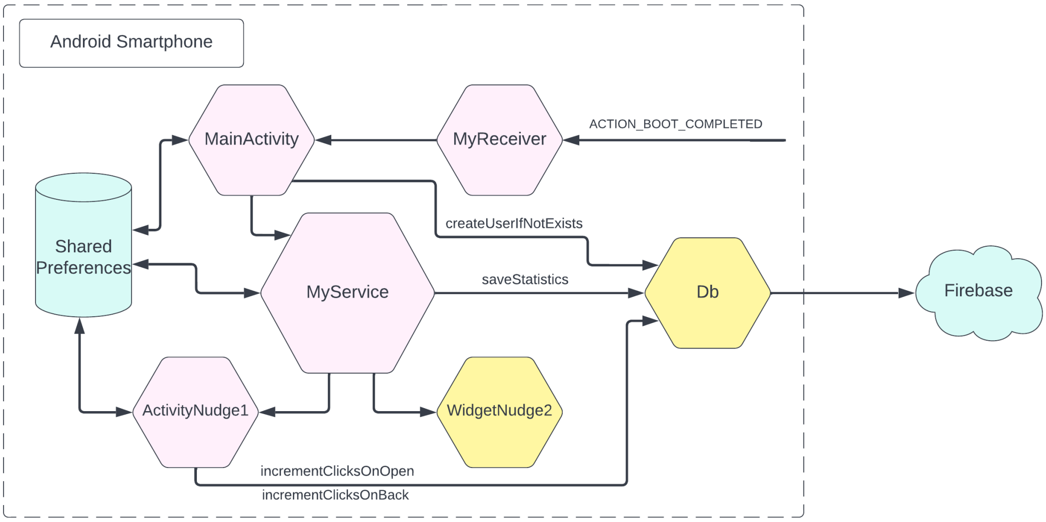

Figure 4.2: Architecture of NudgeApp

perform long-lasting tasks, such as playing music while there is another activity in foreground, or keeping active a VPN connection.

3. **Broadcast receiver**[5], a component that listens and handles special messages sent by the operating system when some events occurs, such as the phone startup or the end of the download of a file. Broadcast receivers cannot have a GUI, but they can generate notifications that are displayed in the status bar.

4. **Content provider**[6], which is a component that manages the application data and provides the standard interface for connecting data in one process with code running in another process.

After this necessary excursus, it is possible to give some further details about the interface architecture.

NudgeApp has been developed as an Android application, using the Kotlin programming language, together with the implementation of some XML files mostly needed for purposes of configuration and for adding resources, and Android Studio as integrated development environment (IDE). The interface, whose structure is depicted in Figure 4.2, relies upon two activities, the first, called ActivityNudge1, for displaying Nudge1 and the other, MainActivity, containing the steps for the

[5]https://developer.android.com/reference/android/content/BroadcastReceiver

[6]https://developer.android.com/guide/topics/providers/content-providers

setup and customization of buttons. Another core component of NudgeApp is the background service, named MyService, which is in charge of detecting wrong behaviors, according to the metrics defined in section 4.2.2, and displaying the two nudges. Nudge2, instead, is modeled as a widget containing an image and showing a textual message when clicked, thus it has not its own activity and exploits the same context of the underlying service. Data collected by the application is saved both locally, using the shared preferences, and remotely, thanks to Firebase[7].

4.3.1 Accessibility service

The major part of the application logic is included in the class MyService, which is an extension of AccessibilityService.

AccessibilityService is a class that models a special kind of background service adopted to "assist users with disabilities in using Android devices and apps"[8], and able to receive callbacks when AccessibilityEvents are fired by the operating system. Examples of these events are the opening and closing of an application, the clicks and swipes on the screen, the taps on buttons, and so on. Accessibility services are really powerful because they are able to interact with events that are raised in external contexts with respect to the one of their "mother" application, i.e., the one which has started the service. For the above reason, their usage is restricted and they need the explicit approval of the user to be run.

Although accessibility services can be configured in code, the best practice is to use an XML configuration file where the properties are set. In the following code snippet they are shown the ones used for this specific case.

```
<accessibility-service
xmlns:android="http://schemas.android.com/apk/res/android"
android:description="@string/accessibility_service_description"
android:accessibilityEventTypes="typeAllMask"
android:accessibilityFlags="flagDefault|flagRequestTouchExplorationMode"
android:accessibilityFeedbackType="feedbackVisual"
android:notificationTimeout="100"
android:canRetrieveWindowContent="true"
android:canRequestTouchExplorationMode="false"
android:settingsActivity="com.example.thesistest2.activity.MainActivity" />
```

[7] https://firebase.google.com/

[8] https://developer.android.com/reference/android/accessibilityservice/AccessibilityService

In particular, it is possible to define a description for the service, to list the AccessibilityEvents which are of interest for the service, to specify the type of feedback provided, to set the timeout, in milliseconds, for receiving events, and to configure other properties such as the possibility of retrieving the content displayed on the screen. Moreover, using the packageNames property, it is possible to define the list of applications for which the service wants to receive the events; in case of NudgeApp, this property was not specified to allow the service to understand better when a social network has been closed, i.e., another app or the default launcher was opened, thus allowing it to hide immediately the Nudge2 widget if it is for some reason displayed on the smartphone screen.

All accessibility service have the possibility of overriding some callbacks used for handling their state, which are available also for standard services. Some examples are *onCreate* and *onServiceConnected*, called before the service is started, while *onUnbind* and *onInterrupt* are called before it is stopped. In case of MyService, the first two callbacks were overridden to load the user settings stored in the SharedPreferences and to create the notification channel and publish a permanent notification, that helps the service in not being considered as "purely background" and, thus, forcely stopped by the operating system. The last two mentioned callbacks, instead, were used to handle the situation in which the service is going to be stopped, by persisting data collected so far both in the local storage and remotely.

Apart from the previous mentioned functions, every accessibility service can override a special callback, named *onAccessibilityEvent*, which receives as only parameter an object modeling the occurred event. Here, it is possible for the service to handle separately the events it is interested in by checking their type and package name. NudgeApp, and the MyService class in particular, handles three main types of events: TYPE_VIEW_SCROLLED, TYPE_WINDOW_STATE_CHANGED, and TYPE_WINDOW_CONTENT_CHANGED.

The TYPE_VIEW_SCROLLED event was used for detecting both scrolls and pulls. When such an event is handled by the onAccessibilityEvent calback, the service gets the length of the scroll, expressed in pixels, by reading the *scrollDeltaY* field of the incoming object. In case the *scrollDeltaY* amount is greater than 0, it means that the user has scrolled down, otherwise he has scrolled up. However, since accessibility events arrive with an interval of 100 milliseconds, which has been set in the service configuration file, at every user interaction may correspond more than one single event. For this reason, in case the amount read is greater than 0 the service starts counting the amount scrolled by increasing a local variable. Finally, when the service receives an event with a length less than a predefined threshold of 200 pixels, it checks the amount scrolled in the local variable and can mark that series of events as a *scroll* and add a record to a local list used for checking the scrolling speed of the user. Having added a threshold of 200 pixels has also been useful for ignoring those small scrolls that are performed not voluntarily when the user is just touching the screen.

Instead, in case of a pull to refresh event, the operating system fires a scrolling event with *scrollDeltaY* set to 0. In this case, the service first checks that also *scrollDeltaX* is set to 0, meaning that the incoming event is not an horizontal scroll, and, for being sure of detecting a real pull gesture, checks also for the presence on screen of a label indicating that the user is at the top of the page, for example the one that invites to publish a post on Facebook. Even in this case, the detected gesture is inserted in a local list of pull events, which is then used to display NUdge2, according to the metrics defined in section 4.2.2.

The TYPE_WINDOW_STATE_CHANGED event was used for detecting when a new social network has been opened. When this happens, the service computes whether Nudge1 has to be displayed or not, according to the metrics discussed in section 4.2.2. If yes, it is created a new Intent, which is the Android mechanism for modeling an operation to be performed and also for starting a new activity. The new intent is then enriched by adding the label and the textual nudges to be shown on screen, and finally it is used to start a new activity which becomes foreground in place of the selected social network. Since the application had to be tested in a user experiment where all participants were native Italian speakers, all labels and, in general, strings used by the interface have been translated both in English and Italian. Among the two possible languages, NudgeApp then chooses the one that fits better with the smartphone environment, depending on the default language set in the system.

TYPE_WINDOW_CONTENT_CHANGED, the last AccessibilityEvent handled by the service, is usually fired by the system whenever the content on screen changes. The management of this event has been particularly useful both to detect when an application has been reopened from the background, and to update the last timestamp of interaction in case the user is seeing a video without touching the display.

Furthermore, this service is also in charge of collecting usage data for the three monitored social networks, i.e., Facebook, Instagram, and TikTok, that are saved both in the SharedPreferences and on Firebase. Specifically, for each of the previously mentioned apps, the accessibility service saves the milliseconds spent inside the platform daily, the number of pull to refresh gestures detected divided according to the levels specified in section 4.2.2, and a list where each record is an array of events considered “important”, i.e., the opening and closing of an app, and the timestamps where Nudge2 for infinite scrolling has been displayed with the relative color. This information is essential to gather some statistics on the users behavior in presence of the interventions, and, thus, on the efficacy of the proposed nudges. Moreover, the service stores the number of times the user clicks on the button for opening the social, and on the one for closing it.

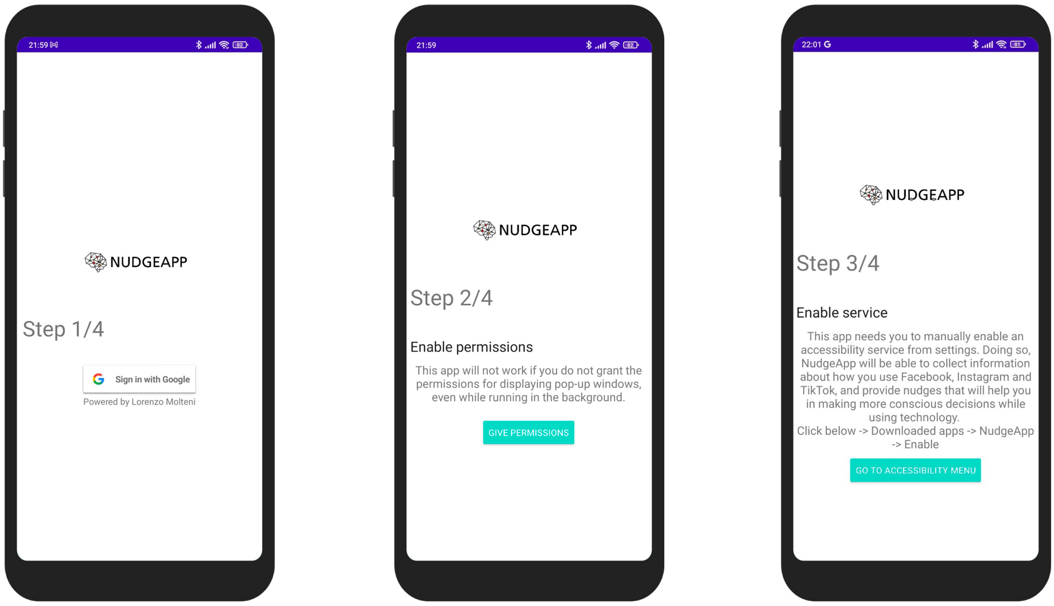

Figure 4.3: Main activity - Steps 1, 2, and 3 of the configuration phase

4.3.2 Main activity

The MainActivity class is an extension of a classic Android Activity, and it is in charge of guiding the user through the necessary phases for correctly installing and running NudgeApp.

This configuration phase is divided in 4 steps (Figures 4.3 and 4.4):

1. The user is requested to login into the application using his Google account. The login process is handled by the One Tap sign-in[9], thus the user just needs to click on the login button and then select one of the possible Google accounts stored in his smartphone from the list that appears at the bottom of the screen. This type of login was particularly useful for two main reasons: on the one side it facilitated the process for storing information remotely, and on the other side it guaranteed a smooth login procedure without the need of managing the registration phase and, thus, storing securely the user personal information and passwords.

2. In step number 2, the user has to give his permissions to the app to display popup windows, even while in the background. These are necessary for NudgeApp to be able to display Nudge1 and Nudge2, since both interventions are started from the background accessibility service.

[9] https://developers.google.com/identity/one-tap/android/overview

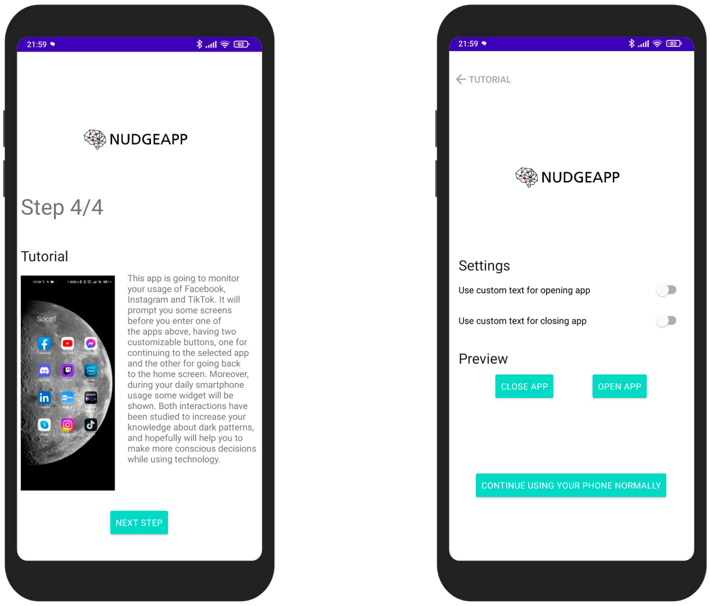

Figure 4.4: Main activity - Step 4 and main page

3. In the following step, NudgeApp requires the user to enable the accessibility service through the specific menu present in the smartphone settings. As previously mentioned in section 4.3.1, accessibility services are really powerful, can see the contents on screen and perform actions in place of users to assist them in their daily life. Thus, this particular kind of service cannot be automatically started in silence, and it needs the explicit permission of the smartphone owner.

4. Lastly, a short tutorial, where the interventions performed by the interface are exposed, is presented to the user. In the same display, NudgeApp provides also a textual description regarding the proposed interventions and the aim of the interface itself.

After this 4 steps configuration, this activity allows the user to see again the tutorial and configure his personal labels that will be displayed in place of the default ones for the open and close buttons shown for Nudge1, i.e., the screen displayed before entering Facebook, Instagram, or TikTok.

4.3.3 Other components

Apart from the main activity and the accessibility service, the project includes other few simpler components, which are listed in the following.

1. ActivityNudge1, which is an extension of an Android Activity, that is started by the background service when the conditions to show Nudge1 are met. It receives, from the intent used to trigger its display, the information needed to populate the screen, including the title, the textual nudge, and a flag indicating whether the nudge to be shown is relative to infinite scrolling or pull to refresh, so that the corresponding gif image is loaded. This activity also provides two buttons, having the default or user-defined labels, the one on the right for opening the selected social network, which invokes the finish() method, and the one left which returns to the smartphone home page before finishing itself.

2. WidgetNudge2, which instead is an object that models the popup widget for intervention number 2. It exposes two functions, the first, called fillWidget, is in charge of setting the proper image and textual nudge, shown when clicking on the popup, inside the widget, while the second, hideWidget, makes the window disappear. Since it is an object, it is not bound to any application or activity, and, thus, it exploits the context of the accessibility service and the frame layout in order to be displayed on screen.

3. MyReceiver, which is an implementation of the BroadcastReceiver, a particular kind of component that is able to listen to events raised by the Android operating system. Specifically, this receiver listens to the ACTION_BOOT_COMPLETED event raised by the OS when the smartphone completes the boot phase and starts the MainActivity, thus giving the user the possibility to manually give again his explicit permission for activating the background accessibility service.

4. Db, which is an utility object created for separating all the implementation related to the remote connection and storage of information on Firebase. It exposes four methods, called by the accessibility service and the two activities, for creating the database record containing the email of the user and for saving the usage data.

Chapter 5

User study

This chapter is dedicated to deepen the goal, the structure and the results of the user study carried out after the end of the interface development.

In this last part of my thesis work, all the participants of the above-mentioned experiment had to download and install properly NudgeApp, and to continue using their phone normally for the whole duration of the experiment while the interface collected their usage data and started displaying the proposed interventions.

5.1 Goal of the study

This final user experiment aims at giving an overall evaluation about the efficacy of the developed smartphone interface, i.e., NudgeApp, whose main goal is to empower users in their daily usage of social networks by giving them the necessary knowledge for taking more conscious decisions, increasing their awareness of malicious designs, and possibly diminishing their technology overuse.

Leveraging the nudge paradigm, which became part of the interface thanks to the interventions planned during the collaborative codesign session described in Chapter 3, NudgeApp tries to inform users about those psychological mechanisms, i.e., dark patterns, that are exploited by social networks to keep everyone online and interactive without his consent. Moreover, the interface wants to make them understand when and how they have been victim of these malicious designs.

In this study, participants had to install the developed interface and to continue using their smartphone as usual for the whole duration of the study, i.e., two weeks, while the application started displaying on their screen the two proposed interventions, i.e., Nudge1 and Nudge2. Apart from the data collected through the accessibility service, as described in section 4.3.1, during this experiment other information, which is further explained in the following section, has been collected. In particular, they include demographic statistics and two self-assessments of users' perceived smartphone addiction, performed at the beginning and at the end of the user study.

All in all, thanks to the variety of information collected, this final user experiment can be seen under two different points of view, which can be summarized in the following research questions:

- **RQ1**: Is the developed interface capable of having an impact in the way the study participants interact with the three selected social networks, i.e., Facebook, Instagram, and TikTok, during their daily life?

- **RQ2**: Have the developed interventions, i.e., Nudge1 and Nudge2, helped participants in reducing their perceived smartphone addiction?

5.2 Structure of the study

The user experiment described in this chapter involved 17 participants, which were invited to download and install properly the developed interface, whose details have been addressed in the previous chapter. All participants were part of my social sphere, including both high school and university students, and young adults. Overall, all the users who have taken part in this study ranged between 18 and 35 years old.

Users were contacted by delivering them a personal WhatsApp text message containing the main information needed to take part in the study. In particular, the message described at high-level the goal of the study itself, the list of the three social networks monitored by NudgeApp, the nature of the data collected by the application and its visibility. Furthermore, this text message was accompanied with some details about the steps needed to succeed in the installation. They included, among the others, a short tutorial showing the installing procedure in a sample phone, and some useful suggestions, dependent on the smartphone manufacturer, to allow the interface to work better, such as blocking the application in the background, allowing it to start automatically and setting the lowest battery consumption restrictions.

More in detail, users were informed that their Google account was needed to run the experiment, that some permissions, like the one for displaying popup windows while in background, were necessary for the app to be effective, and that they had to explicitly give their consent to enable the accessibility service. Even though the whole procedure was driven and assisted by the application, all the users were also informed of the possibility of directly asking for further explanations and details, and for answering any question they may have.

Together with the above mentioned textual message, the participants received through the same WhatsApp chat the application installation package, that is the

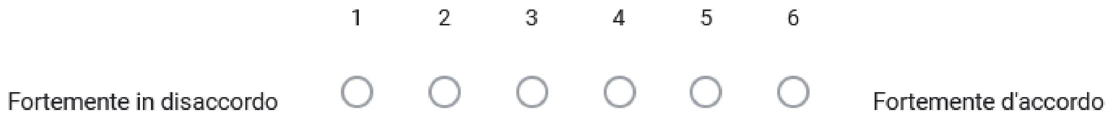

Figure 5.1: Likert scale

file with extension .apk created by Android Studio once the project is built. This package contains all contents of the application, and it just needs to be executed, by clicking on it after having given the consent for installing apps from an unknown source, for completing the installation process.

After having sent their consent to take part in the study, users were requested to fill in two questionnaires created using Google Forms and available online through a link. Google Forms is an online and free to use tool which allows to make questionnaires and surveys giving the possibility of choosing among various kinds of questions, such as multiple choices, linear scale, short answer, paragraph, and checkboxes. Furthermore, it provides graphics, like pie charts and histograms, for easily showing the results.

The first survey included simple questions necessary to collect demographic statistics about age, gender, type of school education, and occupation. In this case, all questions were multiple choices, apart from age and type of school education that were short free answers.

The other one, instead, showed the Italian version of the Smartphone Addiction Scale Short Version [58], thus each of the ten question was a linear scale. SAS [59] is a self-diagnostic scale whose aim is to evaluate smartphone addiction according to six dimensions, which are daily-life disturbance, positive anticipation, withdrawal, cyberspace-oriented relationship, overuse, and tolerance. It was originally developed and validated in South Korea, together with its short version for adolescents, and then it started being validated in other countries, including Italy [60]. The Italian version used in this study includes ten statements, each having six possible answers in a Likert scale (Figure 5.1) from “Strongly disagree” to “Strongly agree”. By analyzing the responses, it is possible to give each questionnaire a score, representing the perceived smartphone addiction of the user.

For both questionnaires, users were informed about the nature of the survey, its estimated duration, and that the collected data would only be used for the purposes of this thesis and only after having been anonymized and aggregated.

The effective user study lasted two weeks, and the application gathered for the whole duration of the experiment the same kind of information about the

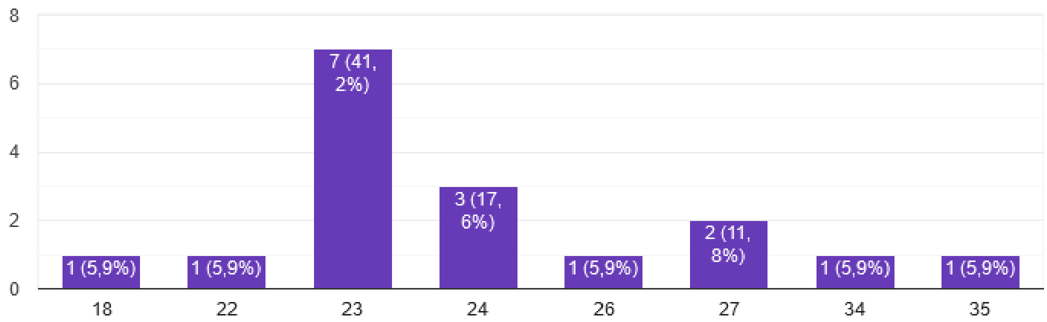

Figure 5.2: Age of user study participants

participants' usage sessions of the three selected social networks, i.e., Facebook, Instagram and TikTok, which have been deeply described in section 4.3.

After the two weeks in which participants tried the proposed interventions, the study ended with a final questionnaire during which users have been asked to answer again to the same online survey used before the beginning of the interface trial, i.e., the Italian version of SAS-SV [60]. Even in this case, users were informed at the beginning of the questionnaire about its goal and duration, and they was told again that all the collected data would have been used only for the purposes of this work, and presented only after having been made anonymous and aggregated.

Collecting the same usage data for the whole duration of the experiment, and letting participants answer twice to the above mentioned self-diagnostic scale, were both useful for accomplishing the dual goal of this user study exposed in the previous section.

5.3 Study results

This section addresses the outcome of the user study, which is the topic of the whole chapter. All results have been retrieved by analyzing the data coming from the three online surveys proposed to users via Google Modules, i.e., demographic information, initial and final questionnaires, and usage data of the monitored social networks, i.e., Facebook, Instagram, and TikTok, which have been collected by the interface and stored remotely using the Firestore Database tools provided by Firebase.

5.3.1 Participants

To run the study, 21 users have been contacted, and 17 of them have correctly completed the user study for its entire duration. Thus, both in this section and in

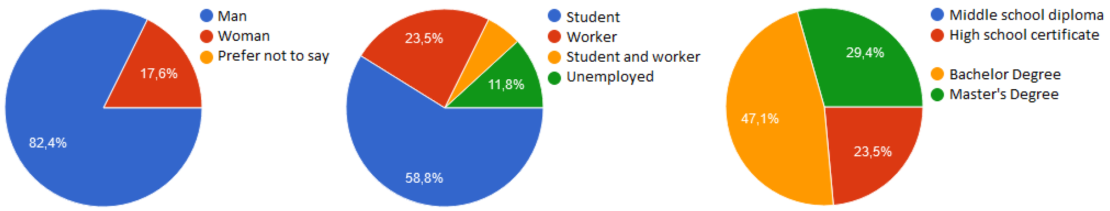

Figure 5.3: Gender(left), occupation (center), and qualification (right) of user study participants

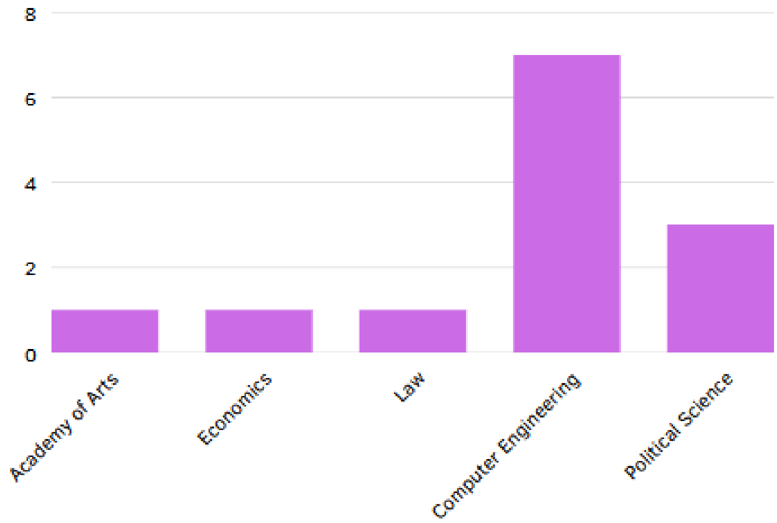

Figure 5.4: Types of school education of user study participants

the following one, graphics and analyses refer to those participants who took part to the whole experiment.

Information exposed in this section was collected thanks to the first survey, i.e., the one concerning demographic information, presented to participants after they have sent their consent to take part to the user study.

From a demographic point of view, as shown in Figure 5.2, participants ranged between the ages of 18 and 35, with 10 of them (58.8%) being 23 or 24 years old; thus the standard deviation is $\sigma = 4.047$. This distribution is perfectly in line with the scope of this thesis, since its goal is to evaluate the impact of nudges on digital wellbeing for the smartphone for teenagers and young adults.

In Figure 5.3, instead, it is possible to have a quick sight on other information:

1. **Gender**: 14 participants out of 17 (82.4%) were men, while 3 were women. None of them decided to not specify his gender.

2. **Occupation**: most of the participants (58.8%) were students, while 4 were workers, two declared to be unemployed and just one to be both, i.e., student and worker at the same time. This is, again, in line with the goal of this study and also with the distribution of ages of the study participants.

3. **Qualification**: 76.5% of the participants said they got a bachelor or master's degree (8 and 5 responses respectively). Just 4 out of 17 were high school graduates, and none of them stopped his studies at secondary school.

Concerning the types of school education and taking into consideration only the ones that have specified to be graduated, thus 13 participants, there is a prevalence of Computer Engineering (53.8%), followed by political sciences (23.1%). Other three academic paths, i.e., Academy of Arts, Economics and Law, were represented by one participant each, as depicted in Figure 5.4.

5.3.2 Efficacy of the proposed interventions

Thanks to the variety of data collected in this user study, it is possible to evaluate the efficacy of the proposed interventions, i.e., Nudge1 and Nudge2, under two different points of view. On the one hand, we can analyze the different answers given by participants to the self-evaluations, performed at the beginning and at the end of the study, in order to understand whether they have changed their perceived smartphone addiction during the experiment and thanks to the implemented nudges. On the other hand, we can look at the usage information collected silently by the accessibility service, as described in section 4.3.1, for understanding whether the proposed interventions have also helped in changing something in the way users behave with their smartphone while they are inside social media platforms.

Self-evaluations

Starting from the self-evaluations, it is first necessary to give some further details about how to give a score to the SAS-SV questionnaire.

Each of the ten questions is represented, as previously mentioned in section 5.2, by a 6 points Likert scale ranging from "Strongly disagree" to "Strongly agree". For scoring each survey there is just the need of assigning to each question points from 1, in case of strongly disagree, to 6, in case of strongly agree, and then to sum all the obtained values.

Depending on the score of the survey, this scale identifies different ranges of addiction for males and females. In particular, males are considered addicted if their score is greater than or equal to 32, while for females this threshold is set to 34 points. Furthermore, the scale also identifies a level associated with high

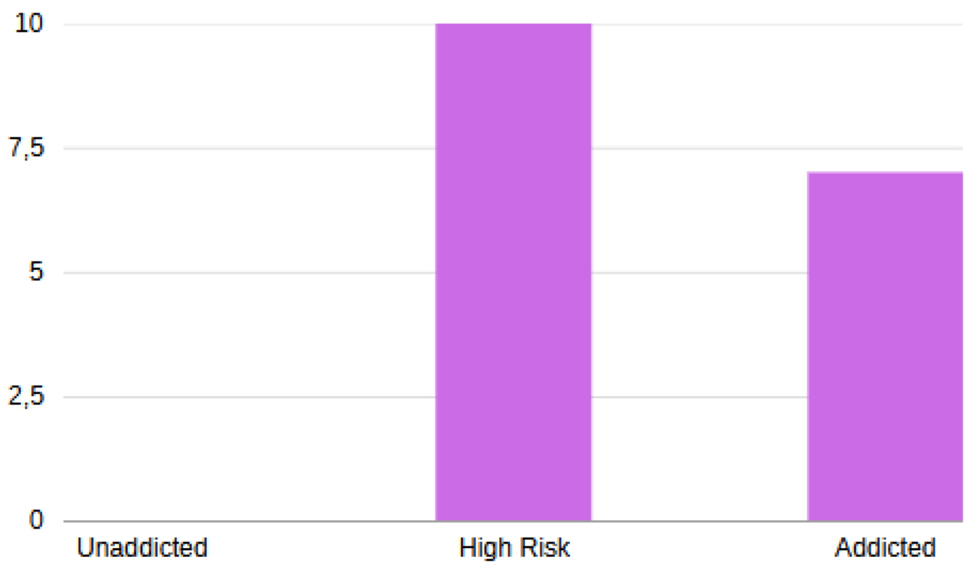

Figure 5.5: Scores of the first SAS-SV questionnaire

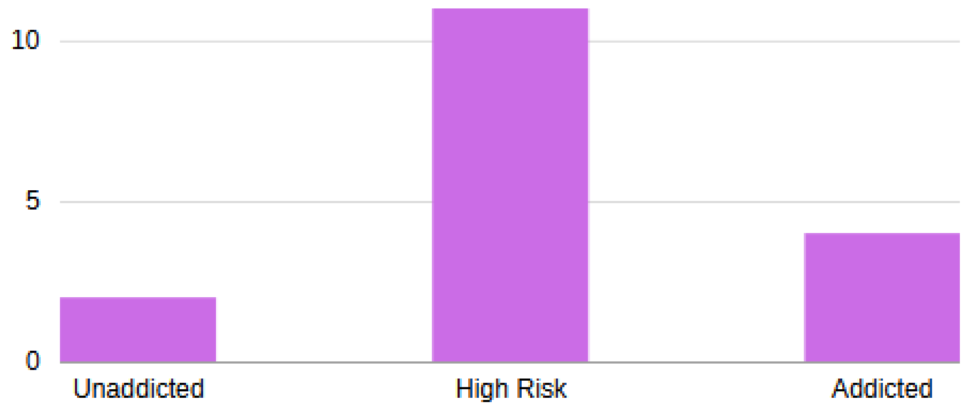

Figure 5.6: Scores of the second SAS-SV questionnaire

risk of addiction, which is represented by scores between 22 and 31 for males, and between 22 and 33 for females.

According to the above-explained levels, and as depicted in Figure 5.5, at the beginning of the experiment 10 participants out of 17 resulted to be at high risk of addiction, while the remaining part, thus the 41.2%, was already considered addicted. Interestingly, none of the participants scored less than 22 points, thus nobody fell into the category of not addicted.

Instead, Figure 5.6 refers to the results of the self-evaluation after the end of the study. Looking at the picture, it is possible to see that the number of unaddicted users increased from 0 to 2, and that the number of addicted participants decreased to 4. The remaining part, thus 11 users, fell into the high risk category. In detail, two participants passed from addicted to high risk, one from high risk to unaddicted, and another one directly from addicted to unaddicted. Even though 5 users increased their score between the first and the second evaluation, none of them have changed his belonging category. Three participants, instead, have scored exactly the same, while all the others, thus 9 participants, which is the 52.9% of the total, have reduced their score.

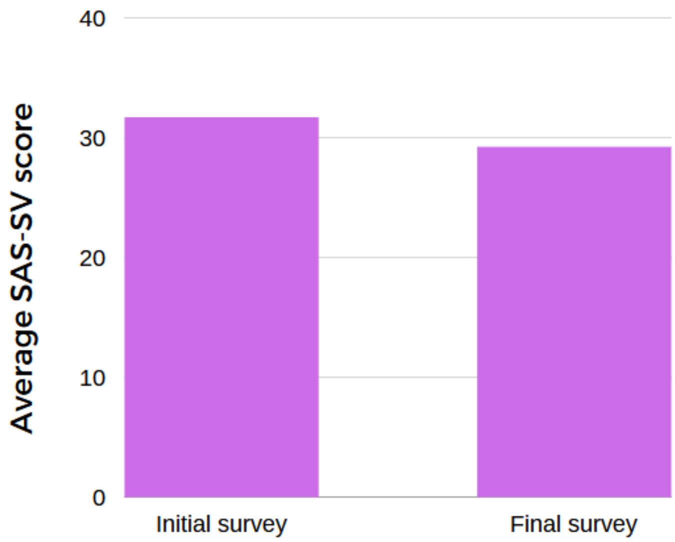

Figure 5.7: Average scores to the SAS-SV questionnaires

Aggregating the results of the performed self-evaluations (Figure 5.7), the average score was 31.7 for the one performed at the beginning (standard deviation $\sigma = 7.044$), while at the end of the study it decreased to 29.2 ($\sigma = 5.631$). This means a reduction of 2.5 points, which, in percentage, means a decrease of around 8%.

Usage data

Concerning the usage data collected during the entire study through the accessibility service of NudgeApp, the first step was to download the entire database content in a format that eases the analysis process. For this reason, and since Firebase does not provide a free to use functionality for exporting data stored in its Firestore Database, I have exploited the *node-firestore-import-export* tool, which is available online under MIT License.

First of all, the tool was installed through the Node Package Manager by running "*npm install -g node-firestore-import-export*". Secondly, it was possible to download all the information stored in the database by running from the shell the command "*firestore-export –accountCredentials credentials.json –backupFile backup.json*".

Using the credentials of the service account associated with the Firebase account, which were created through the project settings in the Firebase console and stored in the *credentials.json* file, this tool has automatically downloaded all the Firestore Database content and saved it into the *backup.json* file. Thus, this implied having a unique but really large file containing the data of all the 17 study participants.

```
import json
f = open('backup.json', "r")
data = json.loads(f.read())
data = data['__collections__']
users = data['users']
for i in users:
    user = users[i]
    #create a file named as the key of the document under the users folder
    with open('users/'+i+'.json', "w") as f:
        json.dump(user, f)
f.close()
```

In order to ease the analysis and to write simpler and smaller programs for extracting relevant information for each single user, I have decided to split the *backup.json* file into smaller JSON files containing just the information of a single individual. For doing so, I have developed and run the code snippet above, written using the Python language and relying on its *json* library. It was used to read the backup.json file produced in the previous step, through the *json.loads* function, to iterate the collection containing the usage data, and to save the statistics of a single user into smaller documents, one per each of the 17 participants of the study, through the *json.dump* method.

Once all the files related to single users have been generated, I have started the actual analysis of the usage data collected by the application during the two weeks experiment. In this phase, the decision was to continue using the Python language, since it provided both an easy and fast way for handling complex objects, such as the one containing the usage information, and a powerful library, *matplotlib*, for printing various types of graphics starting from two simple integer arrays.

For each of the created files, and thus for every user of the experiment, the graphs listed below have been plotted and analyzed:

1. The trend of the minutes spent in total in the three social networks, i.e., Facebook, Instagram, and TikTok.

2. The trend of the amount of pull to refresh gestures performed in total by the user everyday.

3. The trend of the minutes spent daily in total with Nudge2 for infinite scrolling displaying the green, yellow, or red interventions, as described in Table 4.1.

The following snippet, in which the details about how to open and read a JSON file have been voluntarily omitted since they have already been presented previously, represents the Python code used for extracting, analyzing, and plotting the minutes spent in total inside the three monitored platforms as the days vary.

For creating the graph, the *matplotlib* Python library provides the function *plot*, that accepts two arrays of the same size as parameters, representing respectively the values for the abscissa and ordinate axes. The other invoked methods, instead, were used to enrich the graph by adding a title, labels for both the axes, and for setting the position of the ticks on the abscissa axis.

```python
# Conversion from string of values separated by ; to int list
millisFb = user['millisFacebook'].split(';')
millisFb = [int(i) for i in millisFb]
millisInstagram = user['millisInstagram'].split(';')
millisInstagram = [int(i) for i in millisInstagram]
millisTiktok = user['millisTiktok'].split(';')
millisTiktok = [int(i) for i in millisTiktok]
length = len(millisTiktok)
minutesTotal = []
for i in range(0, length) :
    totalMillis = float((millisFb[i] + millisInstagram[i] +
   millisTiktok[i])/60000)
    minutesTotal.append(totalMillis)
# Plot graph
plot.title("Minutes spent in total in the three SNs per day")
plot.xlabel("Day")
plot.ylabel("Time spent (min)")
plot.xticks(range(1, length+1))
plot.plot(range(1, length+1), minutesTotal)
plot.show()
```

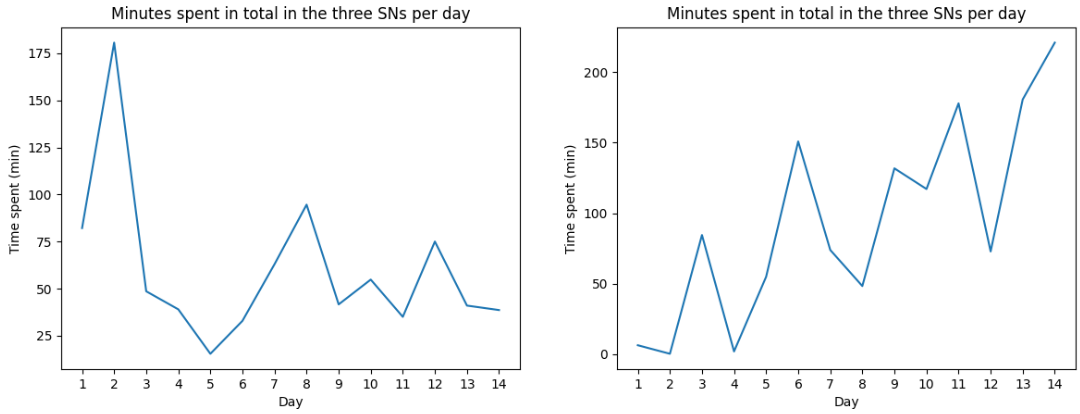

Figure 5.8: Two examples of different trends of total minutes spent on SNs

From the analysis of the first graph, of which two examples are depicted in Figure 5.8, it was not possible to detect any trend common to a significant portion of users. In fact, some of them tended to use the three monitored platforms more in the first week of the study, others increased their usage towards the end of the experiment, and yet others had usage peaks in some days, for example on weekends, probably because of the more free time they had with respect to weekdays.

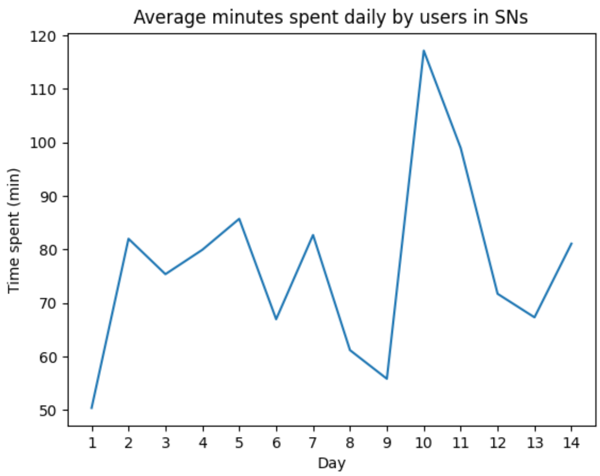

Figure 5.9: Average minutes spent on SNs

Even looking at the average time spent by users during the course of the study (Figure 5.9), it can be seen how, in general, the mean remained almost stable around 80 minutes. In the above graph the curve has its minimum in correspondence of the first day of evaluation. However, it is important to point out that this result could have been driven by the fact that not all users have installed the interface at the beginning of their day, and thus their first usage data could be not complete. Regarding the maximum, instead, it can be seen from the graph how users have spent significantly more time in two consecutive days close to the end of the study, for a total of 117 and 99 minutes respectively. However, since their average usage went immediately back to values similar to the ones before the peak, that is not enough to state that the use of NudgeApp led users spent less or more time inside the monitored social networks.

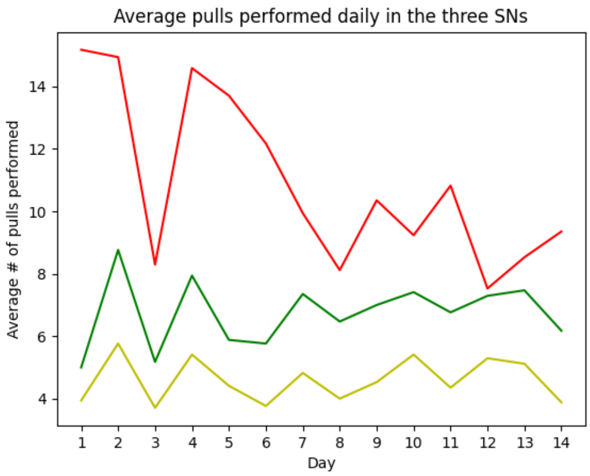

Figure 5.10: Average pulls performed

Regarding the second point, i.e., the trend of the amount of pull to refresh gestures performed daily, it was used a code snippet similar to the one for calculating the minutes spent daily. In this case, it was taken into consideration the number of pulls, but divided by green, yellow, and red, as described in Table 4.2. Thus, the relative graphs contain three different curves. For including more than one line in a single figure using the matplotlib library, it is just necessary to invoke the *plot* function multiple times. Furthermore, it is possible to specify, together with the x-values and y-values of each curve, a third parameter for setting the color of the respective line, e.g., "g" for green, "y" for yellow, "r" for red.

After having calculated the number of pulls for every user, data have been aggregated and the average amount of gestures performed have been computed, always distinguishing between green, yellow, and red interactions. Therefore, the obtained graph, which can be seen in Figure 5.10, shows three different lines, whose colors reflect the distinction mentioned above.

In this graph, it is possible to see how the number of "red pulls" is decreased during the study, passing from an average of about 15.18 at the beginning to 9.35 at the end, which means a reduction of 38.4 percentage points. Day 1 represents also the maximum, while the minimum (7.53) has been reached from this curve in correspondence of day 12, where it is registered a reduction from the highest point of around the 50%. Moreover, this red curve has a mean of 10.91 and a standard deviation of $\sigma = 2.614$.

Instead, concerning the curves for green and yellow pulls, they remained almost stable for the whole duration of the study. In particular, the first had assumed values ranging from around 5 to 8.76, having a mean equal to 6.75 and $\sigma = 1.028$. The yellow curve, instead, ranged between 3.71 and 5.76, with a mean of 4.60 and $\sigma = 0.674$.

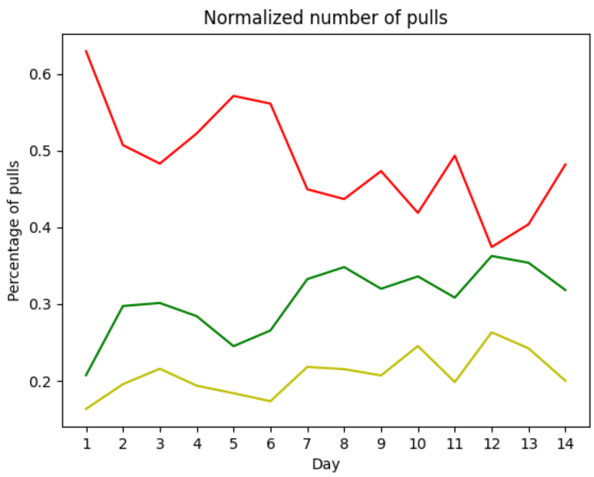

Figure 5.11: Normalized number of pulls

Furthermore, to have a better look on how the behavior of users changed during the study, and to take also into consideration how much they have used their smartphone, the obtained curves have been normalized.

From the original graph, where the lines represented the average gestures performed, another one was generated, where each curve shows the percentage of green, yellow, or red pulls with respect to the daily total. Thus, Figure 5.11 shows the results obtained from the step of normalization, using the usual color scheme. By looking at the picture, we can notice how the number of "yellow pulls" was kept around the 20% of the total for the entire experiment, while green increased from 20% to 31% and red passed from 63% to 48%.

Finally, it has been investigated how users have spent their time scrolling inside the three monitored platforms. As previously exposed in section 4.3.1, the three lists of important events, for Facebook, Instagram, and TikTok, included the timestamps of opening and closing of the app they refer to, and the timestamps corresponding to the display of one of the interventions for infinite scrolling, i.e., the speedometer pointing towards the green, yellow, or red section. Thanks to this information, it has been possible to calculate separately the three amounts of time spent by every user that were considered by NudgeApp, respectively, as normal, fast, and really fast scrolling.

```
for i in range(0, len(impFbEvents)) :
    fbDayEvents = json.loads(impFbEvents[i]['events'])
    for event in fbDayEvents :
        if previousEvent['type'] == 'green':
            greenMillisDay += (event['ts'] - previousEvent['ts'])
        if previousEvent['type'] == 'yellow':
            yellowMillisDay += (event['ts'] - previousEvent['ts'])
        if previousEvent['type'] == 'red' :
            redMillisDay += (event['ts'] - previousEvent['ts'])
        if event['type'] == 'close' :
            if previousEvent['type'] == 'green':
                greenBefore = greenBefore + 1
            if previousEvent['type'] in ['yellow', 'red']:
                otherBefore = otherBefore +1
        previousEvent = event
    greenMillisDay /= 60000
    yellowMillisDay /= 60000
    redMillisDay /= 60000
    greenMinutesTotal.append(greenMillisDay)
    yellowMinutesTotal.append(yellowMillisDay)
    redMinutesTotal.append(redMillisDay)
```

To pursue this goal, a third Python program has been developed, whose crucial parts are reported in the snippet above, that represents the code example in the case of a single social network, i.e., Facebook.

First, the events are retrieved from the user JSON file, converted in a list and ordered by the *date* field, which represents the day they refer to. Then, for every record of the list, that refers to a single day, the inner events are retrieved and iterated for computing the daily milliseconds spent for each of the three levels defined by Nudge2, which are normal, fast, and really fast scrolling, and corresponds, respectively, to the speedometer pointing the green, yellow, or red section. As it can be seen from the code snippet, at every iteration there is a check on the type of the previous event and, depending on its value, the corresponding variable is incremented by the difference between the current event timestamp and that of the previous one.

Furthermore, in case the current event is a *close* one, two local variables are incremented depending on the previous event type. These variables, *greenBefore* and *otherBefore*, are used to count how many times the user was in a "normal scrolling" state before closing a social network, and how many times he was in another, i.e., fast or really fast.

However, it is important to highlight once again that, due to its original length and complexity, the above snippet is just a demo code, where some implementation details, such as control and exit conditions, variable declarations, loop counters increment, and so on, have been voluntarily omitted.

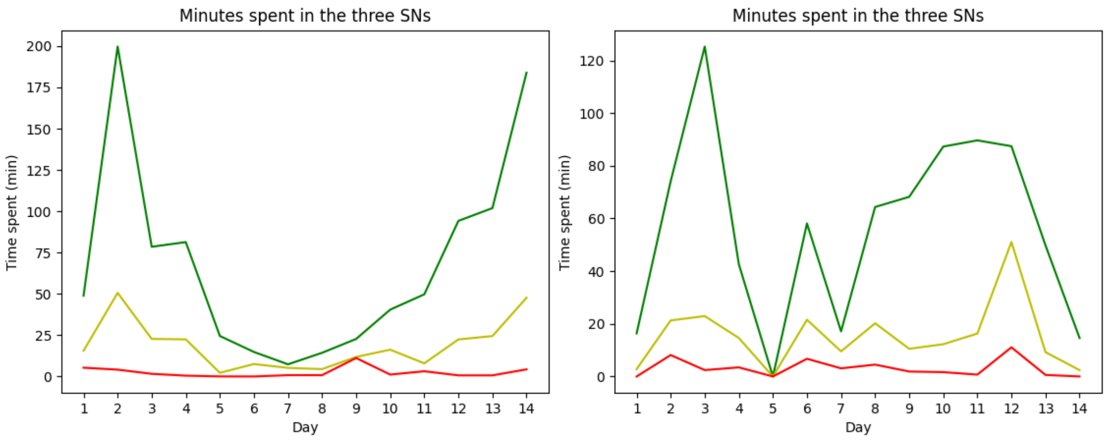

Figure 5.12: Two examples of trends for minutes spent on the three SNs divided in normal, fast, and really fast scrolling

The 17 produced graphs, one per each study participant, of which a couple of examples are depicted in Figure 5.12, represent the individual trend, within the two weeks of the user experiment, of the minutes spent daily adopting a behavior considered by the interface as normal, fast, or really fast scrolling.

From a first analysis of the previously mentioned trends, it was noticed how, in general, the green line lies above the yellow and the red ones, meaning that users have spent most of their time inside the monitored social networks adopting a behavior that was considered "normal" by the interface. In particular, among the data of the 238 days, which is obtained by multiplying 17 participants per 14 usage days each, just in 7 cases, which is around the 3%, the green line was lower than one of the two others, in 29 days (12.2%) the green line was at the same level of another one, or they had a distance which was not visually perceptible from the graph, while in the remaining 84.8% the time spent with a "normal scrolling" behavior was clearly higher than the ones in case of "fast" or "really fast".

Anyway, even if it is true that users have in general spent more time in a "normal scrolling" situation, it is not possible to evaluate the direct impact of the proposed interventions on this behavior. In fact, this distribution of minutes spent could be the same that participants adopt by nature, and could not have been influenced by NudgeApp. Thus, in absence of usage data related to the conduct of users before having installed NudgeApp, that could have been collected, for example, by implementing a control phase at the beginning of the experiment where no interventions are displayed, it is not possible to establish a relation among Nudge2 and the prevalent amount of time spent by participants as "normal scrolling".

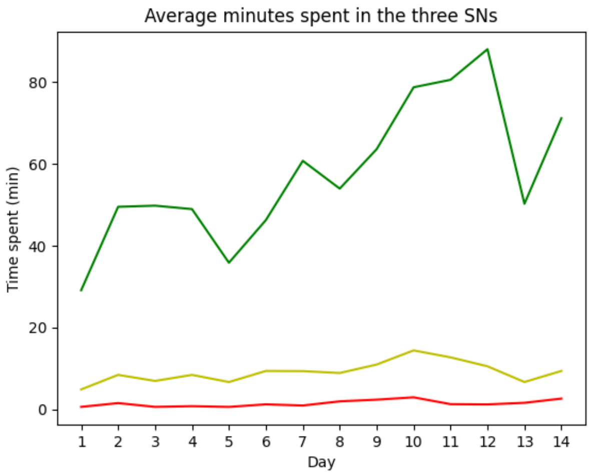

Figure 5.13: Average minutes spent scrolling normal, fast, and really fast

In an effort to better understand the user behavior, the usage data have been aggregated to detect any possible common trend in the way participants have spent their time scrolling inside the three monitored platforms and, most importantly, to recognize eventual changes of their conduct during the course of the study. The results obtained from this analysis are presented in Figure 5.13, from which we can see how, in general, users have scrolled normally, as had already resulted by looking at the 17 individual graphs. Moreover, from the picture it is possible to note the increase in normal scrolling between the beginning and end of the study, that passed from 29.13 to 71.14, reaching its maximum on day 12 (88.0 minutes). However, in the same time interval, even the fast and really fast scrolling curves, i.e., the yellow and red lines, passed respectively from 4.91 and 0.67, to 0.65 and 2.67.

After having normalized these results, thus putting into relation the three obtained curves by calculating their daily percentage with respect to the total, the obtained graph is the one depicted in Figure 5.14. The picture shows how the normal scrolling behavior represented in general around the 84% of the total minutes spent (mean = 84.2%, σ = 1.859), while the remaining was distributed between fast (mean = 13.7%, σ = 1.801) and really fast scrolling (mean = 2.1%,

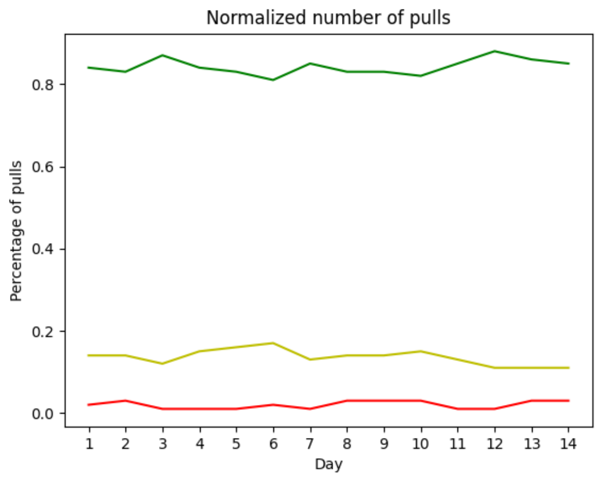

Figure 5.14: Normalized minutes spent scrolling normal, fast, and really fast

$\sigma = 0.925$). However, as it can be seen from the graph, the values have changed little for all the curves, and this is confirmed also from the standard deviations; in particular:

- Normal scrolling ranged between 81% (on day 6) and 88% (day 12).
- Fast scrolling ranged between 17% (on day 6) and 11% (days 12, 13 and 14).
- Really fast scrolling ranged between 1% and 3%.

Therefore, from the above analysis it has not been possible to establish a relation between the usage of NudgeApp and a relevant decrease, or increase, in the time spent on social networks scrolling "normally".

Another interesting consideration can be taken out of the analysis of the widget state before closing one of the platforms, i.e., Facebook, Instagram, and TikTok. Counting the number of times in which the last widget state before a closing event is "green", i.e., Nudge2 is telling the user that he is interacting normally, and the times it is "yellow" or "red", thus representing fast or really fast scrolling, it has been found that the first is higher for *all users*. Diving deeper into the numbers, the first situation has occurred in total 3271 times, while the other just 1212, which are respectively the 73% and 27%. This result is, once again, in line with the proportion between the time spent scrolling normally and the ones scrolling fast or really fast, which were presented in Figure 5.14. Moreover, the 70% of participants, 12 out of 17, have a ratio between the first and the second greater than or equal to 2, meaning that when they close a social network they are frequently, more than 66% of times, in a normal scrolling behavior.

5.4 Discussion

The goal of this work of thesis was to evaluate, through the development and testing of an Android interface, the efficacy of nudges in driving the behavior of users when applied to the digital wellbeing world. The proposed interventions, which have been designed by the six participants of the preliminary collaborative design session, were based on the definition of nudge given by Thaler and Sustein [9] and on the underlying Dual Process Theory [42, 43]. Thus, Nudge1 and Nudge2 tried to alter the user's behavior by showing them a screen before opening one of the monitored social networks or a small widget in the bottom right corner of the display highlighting their compulsive conduct, but left them the possibility of choosing whatever option they liked most, by tapping on the button for opening the social or ignoring the widget on their screen.

After having run the user study described in this chapter, involving 17 participants, and having analyzed the relative results, both in terms of the self-evaluations filled by users and usage data, it is possible to draw some conclusions and to answer to the research question exposed in the introductory chapter.

Regarding the answers to the two SAS-SV questionnaires that all participants were asked to fill at the beginning and at the end of the experiment, it turned out that using NudgeApp has helped them in perceiving themselves as less dependent from technology and from their smartphone. In fact, in the final survey it has been registered a reduction 0f 2.5 points (around the 8%) on the average score obtained by users, meaning that the transfer of knowledge about the dark patterns used in social media from the interface to the minds of users was effective and made them more conscious and less addicted from technology.

Considering the results obtained from the analysis of usage data, it resulted that the proposed nudges decreased the number of "red pulls", i.e., the compulsive usage of the pull to refresh gesture, in favor of the green ones. In detail, the first ones passed from being the 63% of the total to the 48%, while the second increased from 20% to 31%. Thus, we can cautiously state that the proposed intervention for the pull to refresh dark pattern was effective in diminishing the wrong behaviors.

Instead, concerning the average time spent in social networks, and the three different amount of time passed scrolling normally, fast, or really fast, users maintained an almost stable behavior during the whole course of the experiment. In particular, they have spent around 80 minutes per day in total in the three monitored social networks, which is less than the average usage mentioned in the report by Insider Intelligence [57], and around the 84% of this time was considered by the interface as "normal scrolling". Since users haven't changed their conduct during the experiment, it was not possible to establish a relation between NudgeApp and a decrease, or increase, in the above mentioned metrics. However, this result may have been affected by the absence of a control phase at the beginning of the study, which is one of the limitations of this work exposed in Section 6.1, and it might be that the proposed interventions have helped users in controlling their behavior from the first day.

All in all, although a relation between the usage of NudgeApp and a change in the behavior of users was established only for a subset of all the analyzed metrics, this work has proven, at least in part, that nudges are capable of driving the choices of users even while they are using social networks. Therefore, the results obtained by this study have also partially confirmed the outcome of previous researches, in which nudges have been successfully applied in guiding people towards better decisions, both in making healthy food choices [11] and in opting for a pension plan [12].

Chapter 6

Conclusions

At the beginning of this work, some basic information about the underlying world of this thesis and the related work have been presented to the reader. Then, the focus moved on the design and development of an interface, called NudgeApp, willing to increase users' agency and consciousness while using social media platforms, to improve their digital well-being and, hopefully, to diminish their technology overuse. In this step, a codesign process took place, in which 6 habitual smartphone users collaborated to define the interface structure, and its outcome was subsequently used for the development of NudgeApp, an Android application written using the Kotlin programming language that monitors the phone usage and applies the concept of nudge to the digital wellbeing area through two interventions, referred in this work as Nudge1 and Nudge2. Finally, in the last part of this thesis, they are exposed the results of a user experiment, involving 17 participants, whose aim was to evaluate the efficacy of NudgeApp, both in terms of impact in the way users interact with their smartphones while they are inside social media, and in their perceived technology addiction.

The analysis of the above mentioned user experiment, both for perceived technology addiction and social media usage, showed some interesting results. Comparing the self-evaluations about smartphone addiction filled in by all participants at the beginning and end of the experiment, it turned out that NudgeApp helped more than one half of the users to perceive themselves as less addicted, with four of them who passed, according to the ranges defined for the SAS-SV, from a more to a less compulsive situation.

Regarding usage data, participants maintained a stable behavior for the whole duration of the study in terms of minutes spent on the monitored platforms, and it was not possible to establish a relation between the usage of the developed interface and the three amounts of time spent scrolling normally, fast, e really fast. Instead, from the analysis of the aggregated data for the number of pulls performed, divided by green, yellow and red from the less to the more compulsive behavior, it resulted that, thanks to the intervention proposed by interface, the participants reduced, both in proportional and in absolute terms, the number of compulsive interactions, in favor of an increase in the number of gestures considered not problematic.

6.1 Limitations and future work

Although this work showed promising results, it is needed to point out a couple of significant limitations. In the first place, the user experiment performed to evaluate the efficacy of the proposed interventions involved a reduced number of participants (17), with a significant portion of them having a strong technical background (41.2%), and for this reason it is not possible to extend and generalize the obtained results. Secondly, the user experiment only lasted two weeks and did not include a “control phase”, that is a preliminary step where the application to be tested collects usage information in a silent way and without providing any kind of intervention. Thus, a comparison between the users usual behavior and the one obtained after having started using the interface was impossible, making it difficult to understand whether the developed application had also an impact on those metrics that did not change during the study, such as the total time passed inside the monitored social networks and the number of pulls performed. Lastly, this thesis focused only on two selected dark patterns, i.e., infinite scrolling and pull to refresh, and three specific social networks, i.e., Facebook, Instagram, and TikTok, but it is clear that the users digital wellbeing is a way broader field and is influenced by many more factors than the ones taken into consideration in this thesis.

Starting from the results obtained by this work, a first step to consider for future studies is to improve the user experiment by extending its duration, involving a larger number of participants, and including a one-week “control phase” at the beginning, since it is crucial for understanding the current situation of each individual and, thus, for making better comparisons and evaluations on the interface efficacy. Finally, another major improvement could be the design and implementation of other interventions, possibly related to the dark patterns that were not considered in this work, following the same collaborative process used in this study.

Bibliography

[1] Thomas H. Davenport and John C. Beck. «Attention Economy: Understanding the New Currency of Business». In: *Harvard Business School Press* (2001) (cit. on pp. 1, 6).

[2] Alberto Monge Roffarello and Luigi De Russis. «Understanding the Dark Patterns That Steal Our Attention». In: *ACM Trans. Web 00* (2022) (cit. on pp. 1, 7, 8, 16, 19).

[3] Inyeop Kim, Gyuwon Jung, Hayoung Jung, Minsam Ko, and Uichin Lee. «Let's FOCUS: mitigating mobile phone use in college classrooms». In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1.3 (2017), pp. 1–29 (cit. on p. 2).

[4] Jaejeung Kim, Hayoung Jung, Minsam Ko, and Uichin Lee. «GoalKeeper: exploring interaction lockout mechanisms for regulating smartphone use». In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3.1 (2019), pp. 1–29 (cit. on p. 2).

[5] Seekrtech. *Forest - Stay focused, be present.* Retrieved October 2022 from `https://www.forestapp.cc/`. 2022 (cit. on p. 2).

[6] Jaejeung Kim, Joonyoung Park, Hyunsoo Lee, Minsam Ko, and Uichin Lee. «LocknType: Lockout task intervention for discouraging smartphone app use». In: *Proceedings of the 2019 CHI conference on human factors in computing systems*. 2019, pp. 1–12 (cit. on p. 2).

[7] Young-Ho Kim, Jae Ho Jeon, Eun Kyoung Choe, Bongshin Lee, KwonHyun Kim, and Jinwook Seo. «TimeAware: Leveraging framing effects to enhance personal productivity». In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 2016, pp. 272–283 (cit. on p. 2).

[8] Alberto Monge Roffarello and Luigi De Russis. «The Race Towards Digital Wellbeing: Issues and Opportunities». In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. Glasgow, Scotland Uk: Association for Computing Machinery, 2019, pp. 1–14. ISBN:

9781450359702. DOI: `10.1145/3290605.3300616`. URL: `https://doi.org/10.1145/3290605.3300616` (cit. on p. 2).

[9] RH Thaler and CR Sunstein. «Nudge: Improving decisions about health, wealth and happiness». In: *Amsterdam Law Forum*. Vol. 1. HeinOnline. 2008, p. 89 (cit. on pp. 2, 10, 55).

[10] Aditya Kumar Purohit, Louis Barclay, and Adrian Holzer. «Designing for digital detox: Making social media less addictive with digital nudges». In: *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 2020, pp. 1–9 (cit. on pp. 2, 11).

[11] Floor M Kroese, David R Marchiori, and Denise TD De Ridder. «Nudging healthy food choices: a field experiment at the train station». In: *Journal of Public Health* 38.2 (2016), e133–e137 (cit. on pp. 2, 10, 11, 56).

[12] Henrik Cronqvist, Richard H. Thaler, and Frank Yu. «When Nudges Are Forever: Inertia in the Swedish Premium Pension Plan». In: *AEA Papers and Proceedings* 108 (May 2018), pp. 153–58. DOI: `10.1257/pandp.20181096`. URL: `https://www.aeaweb.org/articles?id=10.1257/pandp.20181096` (cit. on pp. 2, 10, 56).

[13] Muin J Khoury and John PA Ioannidis. «Big data meets public health». In: *Science* 346.6213 (2014), pp. 1054–1055 (cit. on p. 6).

[14] Klaus Schwab. *The fourth industrial revolution*. Currency, 2017 (cit. on p. 6).

[15] Kimberly S Young. «Internet addiction: The emergence of a new clinical disorder». In: (2009) (cit. on p. 6).

[16] Klodiana Lanaj, Russell E. Johnson, and Christopher M. Barnes. «Beginning the workday yet already depleted? Consequences of late-night smartphone use and sleep». In: *Organizational Behavior and Human Decision Processes* 124.1 (2014), pp. 11–23. ISSN: 0749-5978. DOI: `https://doi.org/10.1016/j.obhdp.2014.01.001`. URL: `https://www.sciencedirect.com/science/article/pii/S0749597814000089` (cit. on p. 6).

[17] Amy Orben and Andrew K Przybylski. «The association between adolescent well-being and digital technology use». In: *Nature human behaviour* 3.2 (2019), pp. 173–182 (cit. on p. 6).

[18] Jean M Twenge, Gabrielle N Martin, and W Keith Campbell. «Decreases in psychological well-being among American adolescents after 2012 and links to screen time during the rise of smartphone technology.» In: *Emotion* 18.6 (2018), p. 765 (cit. on p. 6).

[19] Saima Salim. *More than six hours of our day is spent online – Digital 2019 reports*. Retrieved November 2022 from `https://www.digitalinformationworld.com/2019/02/internet-users-spend-more-than-a-quarter-of-their-lives-online.html`. Feb. 2019 (cit. on p. 7).

[20] Katie Hurley. *Teenage Cell Phone Addiction: Are You Worried About Your Child?* Retrieved November 2022 from `https://www.psycom.net/cell-phone-internet-addiction`. Oct. 2022 (cit. on p. 7).

[21] Christopher Burr, Mariarosaria Taddeo, and Luciano Floridi. «The ethics of digital well-being: A thematic review». In: *Science and engineering ethics* 26.4 (2020), pp. 2313–2343 (cit. on p. 7).

[22] Dr Paul Marsden. *What is Digital Wellbeing? A List of Definitions*. Retrieved November 2022 from `https://digitalwellbeing.org/what-is-digital-wellbeing-a-list-of-definitions`. July 2020 (cit. on p. 7).

[23] Android. *Digital Wellbeing*. Retrieved November 2022 from `https://www.android.com/digital-wellbeing/` (cit. on p. 7).

[24] Apple. *ScreenTime*. Retrieved November 2022 from `https://support.apple.com/en-us/HT208982` (cit. on p. 7).

[25] Krisztina Szerovay. *Dark Patterns*. Retrieved November 2022 from `https://uxknowledgebase.com/dark-patterns-3b41ed7a690e`. Nov. 2017 (cit. on p. 8).

[26] H. Brignull. *Deceptive Design*. Retrieved April 2022 from `https://www.deceptive.design/` (cit. on p. 7).

[27] Colin M. Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L. Toombs. «The Dark (Patterns) Side of UX Design». In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Available from `https://dl.acm.org/doi/10.1145/3173574.3174108`. Montreal QC, Canada, 2018, pp. 1–14 (cit. on p. 7).

[28] Eric PS Baumer, Rui Sun, and Peter Schaedler. «Departing and returning: sense of agency as an organizing concept for understanding social media non/use transitions». In: *Proceedings of the ACM on Human-Computer Interaction* 2.CSCW (2018), pp. 1–19 (cit. on p. 8).

[29] Justin Cheng, Moira Burke, and Elena Goetz Davis. «Understanding perceptions of problematic Facebook use: When people experience negative life impact and a lack of control». In: *Proceedings of the 2019 CHI conference on human factors in computing systems*. 2019, pp. 1–13 (cit. on p. 8).

[30] Gloria Mark, Yiran Wang, and Melissa Niiya. «Stress and multitasking in everyday college life: an empirical study of online activity». In: *Proceedings of the SIGCHI conference on human factors in computing systems*. 2014, pp. 41–50 (cit. on p. 8).

[31] Alberto Monge Roffarello and Luigi De Russis. «Achieving DigitalWellbeing Through Digital Self-Control Tools: A Systematic Review and Meta-Analysis». In: *ACM Trans. Comput.-Hum. Interact. 00* (2020) (cit. on pp. 9, 10).

[32] Hyunsung Cho, DaEun Choi, Donghwi Kim, Wan Ju Kang, Eun Kyoung Choe, and Sung-Ju Lee. «Reflect, Not Regret: Understanding Regretful Smartphone Use with App Feature-Level Analysis». In: *Proceedings of the ACM on Human-Computer Interaction* 5.CSCW2 (2021), pp. 1–36 (cit. on p. 9).

[33] Geza Kovacs, Zhengxuan Wu, and Michael S Bernstein. «Rotating online behavior change interventions increases effectiveness but also increases attrition». In: *Proceedings of the ACM on Human-Computer Interaction* 2.CSCW (2018), pp. 1–25 (cit. on p. 10).

[34] Geza Kovacs, Drew Mylander Gregory, Zilin Ma, Zhengxuan Wu, Golrokh Emami, Jacob Ray, and Michael S Bernstein. «Conservation of Procrastination: Do Productivity Interventions Save Time Or Just Redistribute It?» In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 2019, pp. 1–12 (cit. on p. 10).

[35] Alberto Monge Roffarello and Luigi De Russis. «Coping with digital wellbeing in a multi-device world». In: *Proceedings of the 2021 CHI conference on human factors in computing systems*. 2021, pp. 1–14 (cit. on p. 10).

[36] Cass R Sunstein. *Why nudge?: The politics of libertarian paternalism*. Yale University Press, 2014 (cit. on p. 10).

[37] Ryan Calo. «Code, nudge, or notice». In: *Iowa L. Rev.* 99 (2013), p. 773 (cit. on p. 10).

[38] Folke Ölander and John Thøgersen. «Informing versus nudging in environmental policy». In: *Journal of Consumer Policy* 37.3 (2014), pp. 341–356 (cit. on p. 10).

[39] Crawford Hollingworth and Lisa Barker. *BE360: Protecting Consumers from 'SLUDGE'*. Retrieved April 2022 from `https://www.research-live.com/article/features/be360-protecting-consumers-from-sludge/id/5031182`. Nov. 2017 (cit. on p. 11).

[40] Verena Zimmermann and Karen Renaud. «The nudge puzzle: matching nudge interventions to cybersecurity decisions». In: *ACM Transactions on Computer-Human Interaction (TOCHI)* 28.1 (2021), pp. 1–45 (cit. on p. 11).

[41] Raymond Samuriwo and Alan D. Pearman. «Bees, owls and cyborgs: The dual process theory & nurses? Pressure ulcer related decision making». In: 2014 (cit. on p. 11).

[42] Magda Osman. «An evaluation of dual-process theories of reasoning». In: *Psychonomic bulletin & review* 11.6 (2004), pp. 988–1010 (cit. on pp. 11, 25, 55).

[43] Daniel Kahneman, Patrick Egan, et al. «Thinking, fast and slow». In: *Farrar, Straus and Giroux* 1 (2011) (cit. on pp. 11, 25, 55).

[44] Ralph Hertwig and Till Grüne-Yanoff. «Nudging and boosting: Steering or empowering good decisions». In: *Perspectives on Psychological Science* 12.6 (2017), pp. 973–986 (cit. on p. 12).

[45] Renata Santiago Walser and Ulrich Remus. «Nudging vs. boosting: designing self-monitoring features for digital wellbeing apps». In: (2021) (cit. on p. 12).

[46] Beyond Sticky Notes. *What is co-design? A brief overview.* Retrieved May 2022 from `https://www.beyondstickynotes.com/what-is-codesign` (cit. on p. 13).

[47] Anna Meroni, Daniela Selloni, and Martina Rossi. *Massive Codesign: A proposal for a collaborative design framework.* FrancoAngeli, 2018 (cit. on pp. 13, 17).

[48] R Edward Freeman and David L Reed. «Stockholders and stakeholders: A new perspective on corporate governance». In: *California management review* 25.3 (1983), pp. 88–106 (cit. on p. 13).

[49] Wikipedia Foundation Inc. *Stakeholder (corporate).* Retrieved May 2022 from `https://en.wikipedia.org/wiki/Stakeholder_(corporate)`. 2022 (cit. on p. 13).

[50] Henri H Achten. «Requirements for collaborative design in architecture». In: (2002) (cit. on p. 13).

[51] Essi Salonen. *A designer's guide to collaboration.* Retrieved May 2022 from `https://designingcollaboration.com/` (cit. on p. 14).

[52] Gary P. Pisano and Roberto Verganti. *Which Kind of Collaboration Is Right for You?* Retrieved May 2022 from `https://hbr.org/2008/12/which-kind-of-collaboration-is-right-for-you`. 2008 (cit. on pp. 14, 15).

[53] Kai Lukoff, Ulrik Lyngs, Himanshu Zade, J Vera Liao, James Choi, Kaiyue Fan, Sean A Munson, and Alexis Hiniker. «How the design of youtube influences user sense of agency». In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems.* 2021, pp. 1–17 (cit. on p. 15).

[54] Justinmind. *The Double Diamond model: what is it and should you use it?* Retrieved May 2022 from `https://www.justinmind.com/blog/double-diamond-model-what-is-should-you-use/`. Dec. 2018 (cit. on pp. 16, 18).

[55] Jonathan Ball. *The Double Diamond: A universally accepted depiction of the design process*. Retrieved May 2022 from `https://www.designcouncil.org.uk/our-work/news-opinion/double-diamond-universally-accepted-depiction-design-process/`. Oct. 2019 (cit. on p. 16).

[56] DYOPATH. *Eight Negative Impacts of Technology*. Retrieved July 2022 from `https://dyopath.com/eight-negative-impacts-of-technology/`. July 2019 (cit. on pp. 26, 27).

[57] Insider Intelligence. *Social Media Update Q1 2021*. Retrieved July 2022 from `https://www.insiderintelligence.com/content/social-media-update-q1-2021`. Mar. 2021 (cit. on pp. 28, 55).

[58] Min Kwon, Dai-Jin Kim, Hyun Cho, and Soo Yang. «The smartphone addiction scale: development and validation of a short version for adolescents». In: *PloS one* 8.12 (2013), e83558 (cit. on p. 40).

[59] Min Kwon, Joon-Yeop Lee, Wang-Youn Won, Jae-Woo Park, Jung-Ah Min, Changtae Hahn, Xinyu Gu, Ji-Hye Choi, and Dai-Jin Kim. «Development and validation of a smartphone addiction scale (SAS)». In: *PloS one* 8.2 (2013), e56936 (cit. on p. 40).

[60] Concetta De Pasquale, Federica Sciacca, and Zira Hichy. «Italian validation of smartphone addiction scale short version for adolescents and young adults (SAS-SV)». In: *Psychology* 8.10 (2017), pp. 1513–1518 (cit. on pp. 40, 41).