23 Ways to Nudge: A Review of Technology-Mediated Nudging in Human-Computer Interaction

Ana Caraban

Instituto Superior Técnico,
Persuasive Tech Lab, Madeira-ITI
ana.caraban@m-iti.org

Daniel Gonçalves

Instituto Superior Técnico,
University of Lisbon

Evangelos Karapanos

Persuasive Tech Lab,
Cyprus University of Technology
evangelos.karapanos@cut.ac.cy

Pedro Campos

Madeira-ITI,
University of Madeira

ABSTRACT

Ten years ago, Thaler and Sunstein introduced the notion of *nudging* to talk about how subtle changes in the ‘choice architecture’ can alter people’s behaviors in predictable ways. This idea was eagerly adopted in HCI and applied in multiple contexts, including health, sustainability and privacy. Despite this, we still lack an understanding of how to design effective technology-mediated nudges. In this paper we present a systematic review of the use of nudging in HCI research with the goal of laying out the design space of technology-mediated nudging –the *why* (i.e., which cognitive biases do nudges combat) and the *how* (i.e., what exact mechanisms do nudges employ to incur behavior change). All in all, we found 23 distinct mechanisms of nudging, grouped in 6 categories, and leveraging 15 different cognitive biases. We present these as a framework for technology-mediated nudging, and discuss the factors shaping nudges’ effectiveness and their ethical implications.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction**.

KEYWORDS

Nudging, Behavioral Economics, Persuasive Technology.

ACM Reference Format:

Ana Caraban, Evangelos Karapanos, Daniel Gonçalves, and Pedro Campos. 2019. 23 Ways to Nudge: A Review of Technology-Mediated Nudging in Human-Computer Interaction. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019)*, May 4–9, 2019, Glasgow, Scotland UK. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3290605.3300733>

1 INTRODUCTION

With the uptake of Personal Informatics (PI) and Behavior Change Technologies (BCT) in research and their diffusion in the commercial world, researchers have increasingly noted the inability of these tools to sustain user engagement. Studies have repeatedly highlighted high abandonment rates [37, 62, 82] and researchers have opted to understand why this happens [61] [26] and how to develop design strategies that sustain user engagement [37], [56], [39].

All this makes sense, as sustained user engagement is crucial to the function of personal informatics as behavior change tools. For instance, research has repeatedly shown that individuals quickly relapse to their old habits once they stop to monitor their behaviors (see [9] for a review). Yet, while finding ways to sustain user engagement is a way forward, and it has received some empirical support (see [37] as an example), researchers have also asked whether such technologies could rely less on users’ will and capacity to engage with the technology and regulate their behaviors (e.g., [1], [63]).

Most PI and BCT tools are, as Lee et al. [63] argue, information-centric. They assume that people lack the knowledge in order to successfully implement changes in their behaviors and the role of the tool is to support people in logging, reviewing and reflecting upon their behaviors. This emphasis on *reflection* as the means to behavior change has recently been noted by Adams et al. [1] who found 94% of behavior change technologies published in HCI to tap to the so-called *reflective mind*, rather than the fast and automatic mental processes that are estimated to guide 95% of our daily decisions.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300733>

Leveraging knowledge from the field of behavioral economics and the concept of *nudging* [88], researchers have designed systems that introduce subtle changes in the way choices and information are presented with the goal of guiding users towards desired choices and behaviors. Yet, while a wealth of “technology-mediated nudges” have been developed and studied over the past ten years, we still have a limited understanding as to how to design effective nudges. In particular, while there is ample discussion on the *why* of nudging (i.e., which cognitive biases can nudges combat), there is very little discussion about the *how* (i.e., what exact mechanisms can nudges employ to incur behavior change).

In this paper we present a systematic review of the use of nudging in HCI research with the goal of laying out its design space. Through an analysis of 71 articles found across 13 prominent HCI venues, this paper makes three contributions to the field. First, it identifies 23 distinct mechanisms of nudging developed within HCI, clustered in 6 overall categories and illustrates how these mechanisms combat or leverage 15 different cognitive biases and heuristics. By doing so, it proposes a framework for the ‘how to’ of nudging that can support researchers and practitioners in the design of technology-mediated nudges. Second, it identifies five prominent reasons for the failure of technology-mediated nudges as discussed in HCI literature. Third, it analyses the ethical risks of the identified nudges by looking at the *mode of thinking they engage* (i.e. automatic vs. reflective) and the *transparency* of the nudge (i.e. if the user can perceive the intentions and means behind the nudge).

2 BACKGROUND

Dual process theories of decision-making

One important contribution to understanding human behavior has been made by Dual Process theories. While differing in their details, they are based on the same underlying concept - we own two modes of thinking: System 1 (the automatic) and System 2 (the reflective) [85][50]. The automatic is the principal mode of thinking. It is responsible for our repeated and skilled actions (e.g., driving) and dominates in contexts that demand quick decisions with minimal effort. It is instinctive, emotional and operates unconsciously. The reflective in turn, makes decisions through a rational process. It is conscious, slow, effortful and goal-oriented (see [50] for a summary of the two modes of thinking).

Both systems cooperate. Yet, as we have a predisposition to reduce effort, the reflective system only comes into action in situations that the automatic system cannot handle [50]. It is estimated that 95% of our daily decisions are not reflected upon, but instead activated by a situational stimulus and handled by the automatic mind [6]. In those circumstances,

we apply *heuristics* - mental shortcuts that enable us to substitute information that is unavailable, or hard to access, for a piece of readily available information that is likely to yield accurate judgments [81]. For instance, when unsure about how to act in a given situation, we may look at what others do and follow their actions, what is known as *herd instinct* [17].

While heuristics support us in making decisions fast and easy, in demanding situations, they also make us susceptible to *cognitive biases* –systematic deviations from rational judgment. For instance, the *status-quo bias* reflects our tendency to resist change and to go along with the path of least resistance [51]. As such, we often chose the default option rather than taking the time to consider the alternatives, even when this is against our best interests. For instance, several countries in Europe have changed their laws to make organ donation the default option. In such so called opt-out contexts, over 90% of the citizens donate their organs; while in opt-in contexts the rate falls down to 15%. Research in the field of Behavioral Economics has provided us with a repertoire of cognitive biases, which can be leveraged in the design of interactive technology that support decision making and behavior change. See Appendix A for a list of the 15 cognitive biases identified in our review, along with examples of interactive technology.

Nudging

Thaler and Sunstein [88] introduced the notion of *nudging* to suggest that our knowledge about those systematic biases in decision making can be leveraged to support people in making optimal decisions. A nudge is defined as “*any aspect of the choice architecture that alters people’s behavior in a predictable way without forbidding any option or significantly changing their economic incentive*”. For instance, changing from an opt-in to an opt-out organ donation policy, as in the example above, has a positive impact on societal welfare, without forbidding individuals’ options or significantly changing their economic incentives. Similarly, replacing cake with fruit in the impulse basket next to the cash register, has been found to lead people in buying more fruit and less cake, when both choices are still available [88].

Over the past 10 years, the idea of nudging has been applied in several domains, including HCI [63][43]. For instance, Harbach et al. [43] redesigned the permissions dialogue of the *Google Play Store* to nudge people to consider the risks entailed in giving permissions to apps, while Lee et al. [63] leveraged knowledge about three cognitive biases to design a robot that promotes healthy snacking. Adams et al. [1] reviewed literature on persuasive technology and classified systems as to whether they focused on the automatic or the reflective mind. They found a remarkable 94% of the systems reviewed to focus on the reflective mind with only

6% (11 out of 176) to focus on the automatic mind. While this result certainly highlights the emphasis persuasive technology places on reflection as a means to behavior change, one should note that nudges do not necessarily tap only to the automatic mind.

Figure 1: Four categories of nudges, adapted from Hansen and Jespersen.

Hansen and Jespersen [42], for instance, distinguish nudges in four categories based on two variables (see Fig. 1): *mode of thinking engaged* (i.e. automatic vs. reflective) and the *transparency* of a nudge (i.e. if the user can perceive the intentions and means behind the nudge). This classifies nudges into ones that intend to *influence behavior* (automatic-transparent, e.g., changing the default option), ones that intend to *prompt reflective choice* (reflective-transparent, e.g. the “look right” painted on the streets of London), ones that intend to *manipulate choice* (reflective-non-transparent, e.g., adding irrelevant alternatives to the set of choices with the goal of increasing the perceived value of certain choices) and ones that intend to *manipulate behavior* (automatic-non-transparent, e.g., rearranging the cafeteria to emphasize healthy items).

3 METHOD

Through a systematic review of novel technological interventions in the HCI field, we aimed at capturing the design space of technology-mediated nudging. Our review followed the PRISMA statement [71], structured in four main phases (see Fig.2).

Identification of potentially relevant entries

Eligibility criteria: All studies published in the top fifteen HCI journals and conferences on Google Scholar Ranking following the publication of the Nudge book (i.e., from 2008 to 2017) were analyzed [88]. For inclusion, articles had to present a novel technology-mediated nudge. Drawing on Hansen [42], we defined nudging as a deliberate change in

the choice architecture with the goal of engineering a particular outcome, and that leverages upon one or more cognitive biases. For this purpose, entries needed to describe the goal of the design strategy and to provide adequate information on the cognitive biases employed. We only considered articles that presented novel prototypes; articles discussing techniques employed by commercial systems were excluded.

Search methods for identification of studies: We focused our search in the top fifteen HCI journals and conferences on Google Scholar Ranking (i.e., CHI, CSCW, Ubicomp, UIST, IEEE TOAC, HRI, IJHCS, TOCHI, BIT, DIS, ICMI, mobile HCI, arXiv HCI, IJHCI and IUI). We excluded arXiv HCI as it does not follow a peer-review procedure. The IEEE Transactions on Affective Computing was also excluded because of weak relevance to HCI and the scope of our review. Instead, we included the proceedings of Persuasive Technology due to its relevance to the scope of our review. Studies were identified by searching electronic databases, scanning references lists of articles and by hand searching. Papers were identified from the following sources: ACM, Springer, IEEE, Taylor and Francis and Elsevier, using the terms “nudge”, “cognitive bias” and “persuasion”. The term “heuristic” was not used as it is employed by different domains to signify different concepts (e.g., performance metrics in Computer Science).

Data collection and analysis: Eligibility assessment for study selection was performed in a standardized manner by the 2 first authors. The search results were logged in an excel file. We removed entries with duplicated titles and entries with blank fields (e.g. missing the title, the keywords or the conference name entirely or partially). The first author carried out this process. To control for interrater effects, the second author performed the same screening on the entire sample. The inter-rater reliability was found to be good (Cohen’s $K = 0.77$). Disagreements were resolved through discussion. After this exclusion, 71 articles remained for the analysis stage. From the 71 articles, 14 papers were added following Wohlin guidelines for snowballing in systematic reviews [97]. The full list of articles can be found in appendix B.

Analysis procedure and dataset description

All in all, a total of 71 papers were selected for further analysis. Most of these came from CHI (32), followed by Persuasive (10), Ubicomp (5) and CSCW (3). We content analyzed all entries based on the exact mechanisms employed and the behavioral economic precepts exploited. The emerging categories were then compared and grouped leading to a total of 6 high-level categories. This procedure was applied recursively until all the nudges clustered in a category shared a number of common attributes (e.g., purpose, heuristic or bias).

Figure 2: Adapted PRISMA flowchart of the article selection process.

Application domains: We identified four prevailing domains: *health promotion* regarding physical activity, smoking, water intake, adherence to medication and others (31%), *encouraging sustainable behaviors* such as recycling, reducing food waste, water conservation or adopting eco-driving conducts (20%), *increasing human performance* such as improving recall or reducing information overload (18%), and *strengthening privacy and security*, such as nudging users away from privacy invasive applications, improving password security and others (9%).

Sample size and study duration: of the 71 papers 55 involved a user study with the proposed system, while 50 of them studied a nudge in an isolated context, thus allowing us to infer their effectiveness. The remaining five studies assessed the combined effects of all nudges and other persuasive techniques used by the system. Only 18 (36%) of the 50 studies had a duration longer than a day and 7 (14%) longer than a month. The majority of the studies were conducted in a laboratory setting. The median sample size was 64 users (min=1, max= 1610).

Technological platforms: Interventions were mostly delivered through web applications (N=27, 38%), physical prototypes (N=14, 20%) such as a key holder that senses the mode of transportation elected (i.e., car or bike) or mobile applications (N=10, 14%). Other solutions involved the use of public displays or smartwatch applications. We could not infer the platforms of 11 entries (16%).

Behavioral Economics precepts: The most frequently used behavioral economics precepts were the availability heuristic (24%), the herd-instinct bias (16%), the status-quo (13%) and the regret aversion bias (11%). All in all, we identified 15 precepts. The full list along with definitions and examples can be found in the AppendixA.

4 RESULTS - NUDGING MECHANISMS

In this section we present our framework of the 23 mechanisms of nudging, clustered in 6 overall categories: *facilitate*, *confront*, *deceive*, *social influence*, *fear*, and *reinforce*. We begin by looking at the motive (i.e., the cognitive biases nudges combat or exploit), and elaborate on the exact mechanisms nudges employ to incur behavior change.

Facilitate

These nudges facilitate decision-making by diminishing individual’s physical or mental effort. They are designed to encourage people to intuitively pursue a predefined set of actions, which resemble people’s best interests and goals. Facilitate nudges exploit the *status-quo bias*, also referred to as the power of inertia, which denotes our tendency to resist change and to go along with the path of least resistance [88], [78]. This predisposition of “*choosing not to choose*” leads us to maintain choices already made because the process of searching for a better alternative is often slow, uncertain or costly [51], [87].

Default options

The power of the default has long been acknowledged to have a significant impact on individuals’ choices. In the digital context, numerous examples can be found. For instance, Egebark and Ekstrom [25] found a 15% reduction in paper consumption when replacing the default printer option to “double-sided print”. *CueR* [4] assigns random, memorable secure passwords to users. The authors found a 100% recall rate within three attempts one week after registration, while significantly improving password security. *DpAid* [98] presents a checklist of symptoms that doctors should consider during diagnosis. Its goal is to mitigate the risk of medical errors by nudging doctors to consider alternative hypotheses and not to stick with the initial diagnosis. *DpAid* was found to lead to a significant increase in correct diagnoses from medical practitioners, and repeated use was correlated with fewer errors [98].

Opt-out policies

Similar to defaults, opt-out policies work by assuming users’ consent to a procedure, leading to automatic enrollment. We found a number of examples of opt-out policies in the technological context. For instance, Lehmann et al. [65] replaced an opt-in policy, where the user was asked to schedule an appointment for vaccinations, with an opt-out policy, where permanent appointments were assigned, assuming prior consent. Participants’ in the opt-out condition were more likely to have an appointment for influenza vaccination, which in turn increased the probability of getting vaccinated. Similarly, *Pixel* [53] attempts to increase password security by

automatically enrolling users to the password generation feature. If a user wants to create her own password, she must opt-out of the feature.

Positioning

Another way to tap into the status-quo bias is by altering the visual arrangement of the options provided. For instance, Turland et al. [92] re-ordered the presentation of wireless networks (i.e. placing the most secure options at the top) and used color codes to label the networks' security (i.e. red color for unsecure networks and green for trusted ones). They found that color and positioning combined led to a significant increase in the rate of secure network selection for 60% of the participants, while nudging by positioning alone was ineffective. Cai et al. [12] rearranged items in a retail website in descending, ascending and random order of product quality. They found that the descending list led consumers to embrace the first option as the reference, serving as comparison for the following items. At the same time, this primed the feature of quality—consumers attributed greater value to the quality of the products as compared to the ascending list, in which consumers attributed greater value to the price of the products. Kammerer and Gerjets [52] observed that individuals rely to a great extent in the rank of search results and changed the interface from list to grid to mitigate this bias. They observed that individuals relied more on source cues rather than the position heuristic to evaluate search results, while eye-tracking data revealed that the majority of users inspected the search results in a nonsystematic way (i.e. free exploration).

Hiding

Similar to the positioning technique, hiding consists of making undesirable options harder to reach. For instance, Lee et al. [63] designed a snack ordering website that aimed at promoting healthy choices. As users browse through a number of pages to find the desired, unhealthy snacks were placed on the last two pages. They found 53% of the participants to opt for a healthy snack.

Suggesting alternatives

Another tactic found in our exploration suggest possible choice alternatives to draw attention to occurrences that might have not been considered. For instance Forwood et al. [33] designed a groceries shopping website to nudge users towards healthier choices. For each food item added to the shopping card, the system search for a possible food alternative (i.e. containing fewer calories within the same food category) and suggest the food swap, either at selection or at checkout. Users were found to accept a median of 4 swaps out of 12 foods purchased, and swaps were more likely to be accepted at selection rather than checkout. Similarly, Forget

et al. proposed *PTP*, a system that suggest more secure alternatives to a user-created password (e.g., changing 'security' to 'Use>curity'). Users can effortlessly review alternative suggestions until they find a memorable one. They authors found this strategy to lead to a significant improvement in the security of users' passwords (i.e. passwords having significantly more estimated bits of security) [32].

Confront

Confront nudges attempt to pause an unwanted action by instilling doubt. Tapping into the *regret aversion bias*—people's tendency to become more careful decision makers when they perceive a certain level of risk [80]—they attempt to break mindless behavior and prompt a reflective choice.

Throttling mindless activity

When battling mindless activity, a simple time buffer to reverse the action can be sufficiently effective. For instance, Wang et al. [94] designed a plugin for the Chrome browser that holds the publication of a Facebook post for 10 seconds, inciting the re-examination of the post's content. Although the countdown could be avoided, their study revealed that several participants reformulated the content and even abandoned the publication during the time interval.

Reminding of the consequences

The *availability heuristic* reflects our tendency to judge the probability of occurrence of an event based on the ease at which it can be recalled [88], [93]. As a result, we might overestimate the probability of events when they are readily available to our cognitive processing (e.g., judging higher than the actual probability of cancer when detecting a lump in our body) while we might be overly optimistic when these events are distant [96]. Nudges in this category prompt individuals to reflect on the consequences of their actions. For instance, Harbach et al. [43] redesigned the permissions dialogue of the Google Play Store to incorporate personalized scenarios that disclosed potential risks from app permissions. If the app required access to one's storage, the system would randomly select images stored on the phone along with the message "*this app can see and delete your photos*". Similarly, Minkus et al. [69] developed a Facebook plugin that confronts the user when disclosing pictures of children: "*It looks like there's a child in the photo you are about to upload. Consider making your account private or Limit the audience of potential viewer*", while Wang et al. designed a web-plugin that aims at mitigating impulsive disclosures on social media through reminding users' of the audience. The system selects five random contacts from the user's friend list, according to the post's privacy setting, and presents the contacts' profile pictures along the message "*These people and [X] more can see this*". The authors observed that while participants did not

feel the need to restrict the privacy of the post, they tended to shape the content as an attempt to eliminate content that could offend others [94].

Creating friction

While remind nudges demand immediate attention and action (e.g., a user is asked to reconfirm her action), friction nudges attempt to minimize this intrusiveness while maintaining the capacity to change users' behavior. Hassenzahl and Laschke [44] talked about the aesthetics of friction and explicated it through a number of prototypes. For instance, *Keymoment* [59] is a key holder that nudges users to choose bike over car by dropping the bike key on the floor when one picks up the car key. *RemindMe* [60] is a wooden ring calendar that allows users to postpone undesirable activities to the future. As the ring rotates clockwise, the time comes and the token drops on the floor. *Forget me not* [60] is a reading lamp that decreases its intensity over time to nudge the user to rethink if it is really needed. Similarly, Moere [70] designed two artifacts that infer the sentiment of chat conversations and use ambient feedback (color and heat) to nudge users to pause and think about what they type. Agapie et al. [3] created an aureole around the query text box, which provides feedback through color and size to motivate individuals to type longer queries in information seeking tasks. The aureole becomes red when the query box is empty or with insufficient information. As the information is added, the aureole starts to fade, becoming blue when the input is perceived as enough to retrieve reliable search results. The authors observed that users typed longer queries in the presence of the halo than in its absence, with an average query length of 6 words.

Providing multiple viewpoints

The *confirmation bias* refers to our tendency to merely seek information that matches our beliefs [73]. This bias leads us to pay little attention to or reject information that contradicts our reasoning. *CompareMed* [66] is a medical decision support tool that collects patients' reviews of medicines from social media and presents two different treatments side by side, while displaying user reviews, thus instigating a comparative inquiry and avoiding a fixation on a single treatment. Similarly, *NewsCube* [75] aims at mitigating this bias by collecting different points of view for an event and offering an unbiased clustered overview. The system collects articles offering different viewpoints, 5 extracts irrelevant data and clusters the information in evenly distributed sections, while identifying the unread sections, to nudge the user to get exposed to all viewpoints.

Deceive

Nudges in this category use deception mechanisms in order to affect how alternatives are perceived, or how activities are experienced, with the goal of promoting particular outcomes.

Adding inferior alternatives

The *decoy effect* refers to our tendency to increase the preference for an option when an inferior alternative (decoy) is added to the original set [8]. For instance, Lee et al. [63] leveraged the decoy effect to promote healthy choices on a snack ordering website. To increase the preference for fruit over a cookie, the picture of a *big* and *shiny* Fuji apple was positioned next to a small withered apple. Adding an inferior, withered apple to the list, the importance of the feature "*shininess*" is increased, leading to the dominance of the shiny apple over all other choices. Similarly, Fasolo et al. [28] motivated the purchase of a laptop in an on-line shopping site by displaying the item next to two other laptops: one of high quality and considerably higher price, and one of lower quality and comparable price.

Biasing the memory of past experiences

The *peak-end rule* suggests that our memory of past experiences is shaped by two moments: their most intense (i.e. peak) and the last episode (i.e. end) [18]. This can have important implications as one could affect how we remember events, for instance, through changing their endings. This would in turn affect future choices, as those are made based on our memory of those events, rather than the actual experience of the events [74], [57]. This idea has been tried in HCI across a number of different contexts. Cockburn explored peak-end effects through manipulating the speed of progress bars [54] and reordering tasks' sequence in a way that the ones demanding lower workload are located in the end [18]. Similarly, Gutwin et al. [41] altered the sequence of events, varying in mental and physical difficulty, in a computer game, and found increased user enjoyment, perceived competence and willingness to replay the game [41], while another tactic induced mistakes by the opponents to boost users' enjoyment of the game at the end of each level [67], [2].

Placebos

The *placebo effect* denotes that the provision of an element that has no effect upon the individual's condition, or his environment, is able to improve her mental or physical response due to its perceived effect [54]. For instance, Beecher [10] noticed that when providing a saline solution to wounded soldiers as a replacement of pain-killing morphine, which was no longer available, soldiers self-reported feeling less pain, although the solution had no physical influence on individuals' condition. We identified two videogames that

used placebos to boost users' self-efficacy and motivation. For instance, the mobile game *Ctrl-Mole-Del*, announces an illusory time extension of 7 seconds to give the opportunity to collect more bonuses. Yet, while the player can collect the bonuses when he hits the target, no reward is provided [23]. Similarly *Wrong Lane Chase*, an arcade racing game, delivers a bonus able to potentially boost the player performance by decreasing the speed of incoming obstacles that should be avoided. However, in fact only the background stage slows down and the obstacles uphold the normal speed. This improved player's performance and diminished players' stress [23].

Deceptive visualizations

The *salience bias* refers to the fact that individuals are more likely to focus on items or information that are more prominent and ignore those that are less so [93]. *Deceptive visualizations* leverage this bias to create optical illusions that alter people's perceptions and judgments. For instance, Adams et al. [1] leveraged the *Delboeuf illusion* to create the mindless plate which attempts to influence individuals' perception of the amount of food that is on the plate. The mindless plate, through a top-down projection, modifies the color of the inner circle of the plate, which causes the portion of food to appear bigger in relation to the spare space on the plate. Hollinworth et al. [46] explored the *Ebbinghaus illusion* and added a circle adjacent to the target, thus making the target appear larger, leading to a significant improvement in the performance of "point and click" tasks among senior users [46], while Colusso et al. [19] adjusted the size of bar graphs to make players think they achieved higher scores than they actually did.

Social Influence

Social influence nudges take advantage of people's desire to conform and comply with what is believed to be expected from them.

Invoking feelings of reciprocity

This approach taps into the *reciprocity bias*, which conveys people's tendency to return with an equivalent action the actions that they received from others [17]. For instance, waitresses that offered a mint along with the bill were found to receive 3% more tips than waitresses that did not offer a gift to their customers [17]. In the digital domain, numerous examples may be found. One of the common examples is found in web platforms, such as the one of Gamberini et al. [35], which offer access to online resources before prompting a request (e.g. contact details). Similar to other social media apps, *Pinteresce*, an interface to Pinterest that aims at reducing social isolation among senior citizens, prompts users to leave comments on others' photo galleries. Due to

the reciprocity effect, users return the action, thus increasing social interaction within the community [11].

Leveraging public commitment

The *commitment bias* is our tendency to "be true to our word" and keep commitments we have made, even if there is evidence that this is not paying off [84]. For instance, getting people to verbally repeat a scheduled appointment with their doctor prompts decisions in consistency with the agreement made [42]. Cheng et al. [15] leveraged the commitment bias to reduce the risks of student drop-outs in large online classes. They added a simple button located at the top of the assignment webpage with the message "I've started on this Assignment". When clicked, the button turns green, and the system logs the student's progress through the assignment, and shares this with the course's instructor. This approach was found to instigate higher task compliance and goal achievement [15].

Raising the visibility of users' actions

This approach leverages the *spotlight effect*, our tendency to overestimate the extent to which our actions and decisions are noticeable to others [36], thus promoting behaviors that elicit social approval and avoid social rejection. For instance, electronic boards that make one's real-time speed public, nudge users to adjust their speed and comply to the norms [42]. Examples of digital nudges include *SocialToothbrush* [13] which tracks users' brushing frequency and performance and provides persistent feedback through light. Thus, a parent may quickly notice when entering the bathroom that the child has not (adequately) brushed her teeth. This form of *social translucence* [27] is assumed to nudge the child towards the desirable behavior through increasing *mutual awareness* (i.e., the child knows that her parent will know). *BinCam* [90] a system that integrates a smart phone on the underside of the bin's lid, which captures the waste produced by a household every time the phone's accelerometer senses movement. The photos are automatically shared to all *BinCam* members on Facebook and the owners can see who recently viewed the photos.

Enabling social comparisons

The *herd instinct bias* refers to our tendency to replicate others' actions, even if this implies overriding own beliefs [17]. According to Festinger [30], we tend to pay attention to others' conducts and search for social proof when we are unable to determine the appropriate conduct. Our interactions with quantified-self technologies are filled with such moments. As Gouveia et al. [38] suggest, even the seemingly simple display of Fitbit Flex, with five LEDs that illuminate for each 20% of a daily walking goal achieved, requires some quite difficult projections, if one wants to use it for immediate

self-regulation: “*If I have walked 4000 steps by noon, is this good enough? Am I likely to meet my daily goal?*”. Rather, one might enable a direct comparison to others’ behaviors: “*Have I walked more than what I had done yesterday at the same time? Have I walked more than others having the same daily step goal?*”. Eckles et al. [24] explored different persuasive strategies through mobile messaging and found that the message “*98% of other participants fully answered this question*” led to a significant increase in the disclosure of the requested information. Selecting appropriate comparisons is critical to the success of the nudge. For instance, Colusso et al. [19] found that comparing game players’ performance that exhibit similar levels leads to higher game performance. Similarly, Gouveia et al. [38] developed *Normly*, a watch face that continuously visualizes one’s walked distance and that of another user that shares the same daily step goal, through two progress bars that advance clockwise. They found that when users checked their watch and were not far ahead or behind others, they were more likely to initiate a new walking activity in the next 5 minutes, as compared to the times when the distance between the two users was higher.

Fear

Fear nudges evoke feelings of fear, loss and uncertainty to make the user pursue an activity.

Make resources scarce

One approach is to reduce the perceived availability of an alternative in terms of quantity, rarity or time. The *scarcity bias* refers to our tendency to attribute more value to an object because we believe it will be more difficult to acquire in the future [17]. For instance, announcing limited seats at future events increases the probability of people committing to attend the event well in advance [17]. Cialdini [17] has theorized that the fear of missing out on the opportunity (loss aversion), drives people to action not for the real need of the object, but for the need to avoid the feeling of a loss. Kaptein et al. [55] leveraged the scarcity bias through persuasive messages such as: “*There is only one chance a day to reduce snacking. Take that chance today*”. Gouveia et al. [38] designed *TickTock*, a smartwatch interface that displayed one’s physical activity of only the past hour, thus making feedback a scarce resource. They found that this strategy led users to check their watch more frequently (i.e. on average every 9 mins) and led to a significant increase in physical activity.

Reducing the distance

We often fail to engage in self-beneficial activities when the outcomes are distant in time (e.g., saving for retirement), or

hypothetical (e.g., buying a smoke alarm). Nudges in this category act through reducing these forms of psychological distance [91]. Zaalberg and Midden [99] found that simulating flooding experiences (e.g. listening heavy rainfall and observing the river to rise slowly) was able to motivate individuals to acquire a flood insurance, while Chittaro [16] explored gain versus loss framed messages to encourage people to acquire smoke alarms. They found gain-framed messages (e.g., “*With smoke alarms, you are warned that smoke is entering your bedroom while you are sleeping: in this way, you wake up in time, when it is still possible to escape from the building*”) to be more effective with women, while loss-framed messages (e.g., “*The causes of a fire can be very common and trivial: a short circuit in an electrical appliance, a cigarette left burning on an ashtray, a cloth over a lit lamp. The absence of a smoke alarm does not allow us to detect these events early and imprisons us in a deadly trap*”) to work better with men. Gunaratne and Nov [40] leveraged the endowment effect, our tendency to overemphasize the value of objects we own [51], to design a system that supports users in selecting a retirement savings plan. The interface sets a savings goal and allows the user to observe the predicted outcomes of different retirement plans. The system displays the discrepancy between the goal and the expected saving, emphasized in red, thus the goal is perceived as an endowment, which influences users to adjust their savings decision to preserve the endowment.

Reinforce

Nudges in this category attempt to reinforce behaviors through increasing their presence in individuals’ thinking.

Just-in-time prompts

Just-in-time prompts draw users’ attention to a behavior at appropriate times (e.g., when a behavior deviates from the ideal). For instance, *WalkMinder* [45] buzzes when the user is inactive for prolonged periods, while *EcoMeal* [58] weights the food on one’s plate, infers the eating pace and nudges the user to slow down through light feedback. Similarly, *Eco-driving* [64] provides light feedback when the driver deviates from fuel-efficient driving, while the Smart Steering Wheel [47] vibrates when aggressive driving conducts are inferred. McGee-Lennon et al. [68] explored the use of auditory icons to support medication adherence among senior citizens, while Zhu et al. [100] used pop-up reminders to encourage posture correction when working at the computer.

Ambient feedback

Ambient feedback attempts to reinforce particular behaviors while reducing the potential disruption on users’ activity. For instance, Rogers et al. [79] used twinkling lights to reveal the path towards the closest staircase, while *Gurgle* [5] is a water fountain installation that emulates a rippling water

illusion in the presence of a passerby to motivate water intake. Similarly, Jafarinaimi et al. [48] and Ferreira et al. [29] presented interactive sculptures that mimic office workers' posture in an attempt to break prolonged sedentary activity.

Instigating empathy

Tapping on the *affect heuristic* which denotes that, given that our first responses to stimuli are affective, they have a strong influence on decision making [83], empathy nudges leverage emotionally charged representations to provoke feelings of compassion. One example is the *Never Hungry Caterpillar* [60], an energy monitoring system that uses the representation of a living animal, a 'caterpillar', to display feedback and engage users in sustainable behaviors. When the system detects 'ideal' energy consumption, the 'caterpillar' extension breathes gently and slowly. When behaviors deviate from the ideal (e.g. leaving a device on stand-by mode), the extension starts twisting in pain. Similarly, Dillahunty et al. [21] studied the efficacy of different emotionally engaging visualizations, such as sunny versus dark and stormy environments, or a polar bear whose life is threatened, to motivate pro- environmental behaviors among children. Lastly, *Powerbar* [20] attempts to motivate eco-friendly behaviors through enabling users to donate the savings to childhood care and education institutions, depicting information about the receiver, the location and the purpose.

Subliminal priming

A different way to reinforce a behavior is through the priming of behavioral concepts subliminally, that is below levels of consciousness [86]. While the subliminally presented stimuli do not affect individuals reasoning, they can trigger action by making the representation of the behavior available in the unconscious mind. This is assumed to happen due to the *mere exposure effect*, which conveys that the prolonged exposure of a stimulus is sufficient for increasing a predisposition and preference towards it [76] quickly flashed goal-related words of physical activity (e.g. active) to unconsciously activate behavioral goals each time the user unlocked his phone. Caraban et al. [14] developed *Subly*, a web-plugin that manipulates the opacity of selected words as people surf on the web, while Barral et al. [7] quickly flashed certain cues to encourage particular food selections in a virtual kitchen.

5 DISCUSSION

All in all, we found 74 examples of nudging in HCI literature. Our analysis identified 23 distinct mechanisms of nudging, grouped in 6 overall categories, and leveraging 15 different cognitive biases.

A frequent criticism to nudging is that it works through the manipulation of people's choices. One should note that this is not necessarily the case. Building on Hansen's and

Jesperen's [42] taxonomy, we attempted to position all 23 mechanisms along the two axes: *mode of thinking engaged* (i.e. automatic vs. reflective) and the *transparency* of a nudge (i.e. if the user can perceive the intentions and means behind the nudge; see Fig. 3).

The majority of the examples reviewed (N=39, 52%) were nudges that work through *prompting reflective choice* (top-right quadrant). An example of this is Minkus' et al. [69] Facebook plugin that confronts the user when disclosing pictures of children and suggests making one's account private or limiting the audience of potential viewers. The second most frequent type of nudges found were interventions that *influenced behavior* (N=19, 26% bottom-right quadrant). While tapping on the automatic mind, these nudges are transparent regarding their intentions and means of behavior change. As such, individuals are provided with the option to ignore or reject the suggested choice. An example of this is Rogers' et al. [79] twinkling lights that reveal the path towards the closest staircase, where individuals can interpret its intentions and act accordingly.

Figure 3: Nudges positioned along the transparency and reflective-automatic axes.

Last, we found 16 examples (22%) that work through *manipulating behavior* (bottom left quadrant). Such nudges may raise ethical questions as their intentions and effects are likely to go unrecognized by users. For instance, in the case of opt-out policies users may not recognize their automatic enrollment in a procedure. However, one may argue that not all opt-out mechanisms raise ethical questions. For instance, Lehmann et al. [65] found that their automatic enrollment into the vaccination program was cancelled by 61% of the users, which implies that it is unlikely to have gone unnoticed and users maintained their freedom of choice. The visibility of the automatic enrollment and the ease with which users may opt-out are factors that will influence the ethics of this nudge. Overall, this analysis highlights that there is no latent

conflict between nudging and transparency, nor one between nudging and reflective thinking.

When do nudges fail?

49 (66%) of the 74 nudges that were empirically studied were found to have a significant effect on target behaviors or attitudes. We found no overt relation between the exact nudging mechanism and its effectiveness. Rather, the effectiveness of the nudge depended on its particular implementation in a given context. In this section, we discuss the main reasons of failure, as identified in the articles we analyzed.

Lack of educational effects

One possible risk in using nudging, and especially techniques that tap onto the automatic mind, is their lack of any educational effects [63]. One may then wonder up to what extent the effects of nudges persist when those are removed. For instance, Egebark and Ekstrom [25] found that while the effect of the default “double-sided print” option sustained even after six months, it quickly disappeared when users started using new printers with single-sided print as the default option. On the contrary, Rogers et al. [79] found the effect of twinkling lights that reveal the path towards the closest staircase, to sustain over 8-weeks, and most importantly, even when the nudge was accidentally removed due to a wiring problem. Contrary to the first case, the twinkling lights, while being subtle, might have invited users, in some cases, to reflect over their behaviors, which might have led to the *introjection* of using the stairs as a self-beneficial activity.

Nudging effects not sustaining over time

Besides the lack of educational effects, other factors may also influence the sustainability of nudging effects. For instance, identifying a placebo after repeated exposure might alienate users and provoke feelings of distrust on the system [23]; the effects of subliminal cueing might degrade over time [7]; reminders might cause friction and reactance after repeated exposure [89] and graphic warnings can lose their resonance over time [89]. It is important to acknowledge that while some of the nudges may be more effective at the initial acquisition of a behavior, others might be better supporting behavior maintenance.

Quite surprisingly, only 18 (35%) of the 52 studies had a duration longer than a day, and 10 (19%) one longer than a month. Similarly, only 7 (14%) inquired into whether the effects sustained after the removal of the nudge. This suggests that while initial results are promising, we have a very limited understanding of the long-term effects of nudging in a technological context. Future studies should invest on field-trials of technology-mediated nudges to inquire into their effects over the long term and once nudges are removed.

Unexpected effects and backfiring

Nudges may also backfire and produce unexpected effects, due to *compensating behaviors* (e.g., printing more when double-sided due to carrying less weight, increasing calorie intake along with physical activity due to a licensing effect), *unexpected interpretations* (e.g., showing the average household energy consumption has led individuals to consume more when they notice that they are below average) and other reasons. For instance, in the study of *Pixel* [53], the authors found that because users knew that they would receive a new auto generated password soon, they avoided the extra step to create their own password. Munson et al. [72] found that asking people to make their commitments public led them to make fewer commitments, as people feared the possibility of criticism. Gouveia et al [38], found that enabling social comparisons increased users’ motivation only when their performance was similar to that of others. This implies that the nudge might even have a negative effect when this condition was not met. All in all, we found the majority of the studies not to inquire into possible backfires and unexpected effects. We believe that a stronger emphasis is required in order to advance our understanding of the conditions for effective nudging.

Intrusiveness and reactance

Some nudges work through creating friction. These nudges run the risk of reactance as they thwart people’s autonomy. For instance, Wang et al. [94] reported that at least one of the participants felt censored by the time buffer added to her Fb posts, and many more found it frustrating and time consuming. Similarly, Lehmann et al. [65] found people to unroll from the vaccination program, likely because they felt their autonomy was taken away. Laschke and Hassenzahl [60] presented a design exploration into the *aesthetics of friction* —a set of principles for products that create friction but not reactance.

Timing and strength of nudges

Finetuning the timing and the strength of nudges can be of paramount importance. For instance, Räisänen et al. [77] displayed warning pictures related to the dangers of smoking and observed that the opportune moment to show these pictures, was not when individuals were already smoking but rather much earlier. Forwood et al. [18] explored peak-end effects in the context of HCI, and attributed insignificant results to a weak manipulation of the ending experiences. Similarly, Lehmann et al. [65] found that making the opt-out of an auto-enrollment program too easy (i.e. simply following a link in the email), led to a substantial opt-out rate.

Strong preferences and established habits

Nudges work best under uncertainty that is when individuals do not have strong preferences for particular choices, or established habits. For instance, Forwood et al. [33], who created a system that nudges individuals towards healthier food choices, suggested that a number of factors such as the strength of preferences for certain food choices, and the extent to which food choices are habitual, can influence the effectiveness of the nudge. Räsänen et al. [77] observed that the less a user smoked, the more affected he was by the smoking cessation nudges. Similarly, commitment nudges to enroll into a vaccination program were not effective for individuals with strong negative attitudes towards vaccination [65]. This ineffectiveness of nudges, however, as Sunstein [89] suggests, should be seen in a positive light. If a nudge does not work for particular people and in particular contexts, it may imply that the nudge preserves individuals' freedom of choice – “if choosers ignore or reject it, it is because they know best” ([89], p. 3). At the same time, we believe that there is an untapped opportunity for the personalization and tailoring of nudges. While nudging was initially conceived as a one-size-fits-all approach, technology provides new opportunities as nudges can be tailored to particular contexts; some of us may be more susceptible to particular nudges than others, and some nudges may be more effective in particular contexts than others.

Nudges as facilitators, sparks or signals

Which types of nudges should we use in different situations? We attempt to provide a first answer to this question by mapping the 23 mechanisms into the types of triggers as suggest by Fogg's Behavior Model [31]. This model suggests that for a target behavior to happen, a person must have sufficient motivation, sufficient ability, and an effective trigger. It identifies three types of triggers: *facilitators* (ones that increase ability to pursue the behavior), *sparks* (ones that increase motivation for the pursue of the behavior, and *signals* (ones that merely indicate or remind of the behavior). To understand the function and strengths of the 23 types of nudges, we classified them in each of the triggers (see Fig. 4).

Facilitator nudges aim to simplify the task and make the behavior easier. They are suitable in situations where the user has the motivation to perform the behavior but lacks the ability to do so, such as when there are too many options available, or when the user lacks the ability to discern between the different options. Take as an example the case of defaults and opt-out mechanisms: nudges can increase ability by reducing the physical or cognitive effort required for a course of action. Additionally, facilitators can also be designed to battle impulses by putting additional effort into choosing or by prompting reflective choice (i.e., suggesting alternatives).

Spark nudges are suitable in situations where the user has the ability but it is not sufficiently motivated to pursue the behavior. Sparks are designed to include one or more motivational elements. We observed that spark nudges increase motivation through leveraging the perceived likelihood of a loss; by increasing individuals' self-efficacy (e.g., through the use of placebos); by supporting planning (e.g., public commitment); by increasing accountability and personal control (e.g., increase visibility of users' actions); by inserting competing attractive alternatives (e.g., decoy and deceptive visualizations); or, by exploiting social acceptance mechanisms (e.g., social comparisons).

Signal nudges are suitable in situations when both motivation and ability are present but there is a discrepancy between users' intentions and their actions. Nudges, in this case, reinforce behaviors either by triggering doubt (e.g., through friction), by triggering discomfort with the current behavior (e.g., instigate empathy), or by increasing the preference to certain stimuli (e.g., subliminal priming). In these cases, users can avoid aversive tasks, stop the action to avoid the tension caused or can engage with a behavior that is in the top of their mind.

Figure 4: The 23 nudging mechanisms mapped into three types of triggers suggest by Fogg's Behavior Model.

Design considerations

What should designers consider when designing new types of nudges? Table 1 outlines design considerations for three of the 23 nudging mechanisms (for the full table consult Appendix C).

For instance, in the case of “suggesting alternatives” one may ask how many alternatives should be suggested, when should they be proposed, and how should they be presented

Table 1: Design considerations for five out of the 23 mechanisms. See Appendix C for the full table.

<i>Mechanism</i>	<i>Design considerations</i>
Suggesting alternatives	How many alternatives should the system suggest? When should they be presented (e.g., during, before or after a selection has been made)? What type of suggestions should be made? <i>Hint</i> : It is important to determine the number of choice alternatives and attributes users can process without suffering the negative effects of overload.
Default options	What constitutes an appropriate default choice or value, and why? Should the default be personalized or adapt over time (e.g., gradually reducing the size of a plate in a restaurant)? Who bears the ethical responsibility when an inappropriate default is presented and unwanted consequences arise, for instance, in the case of algorithmic decisions.
Reminding of the consequences	What are the main undesirable consequences of the behavior to be altered? Are they severe enough to dissuade the behavior when presented by the system? How can you alter users' perception of the likelihood of their occurrence? How can the system make the consequences, in terms of losses, more personal?
Placebo	What is the primary function of the placebo (e.g., to increase self-efficacy?). How this can be achieved? How can the system make the user feel in control? Can you ensure that the information presented is noticeable, yet trustable?
Make resources scarce	How can the system render the desirable alternative as a scarce resource and invoke feelings of missing out if not pursued? Is the use of text, images or visualizations more appropriate? <i>Hint</i> : Using language, which implies that the audience already has achieved the outcome or selected the alternative, can trigger feelings of ownership and in turn, increase users' motivation to avoid a loss.

to the user (e.g., whether one should be visually highlighted or preselected)? Similarly, in the “*default*” mechanism, designers should think about what constitutes an appropriate default choice or value and why; how easily can users opt-out of the default and what effect will this have on users' autonomy and the efficacy of the system; or, should the default be tailored to each individual and who bears the ethical responsibility when an inappropriate default is presented and unwanted consequences arise?

Similarly, when designing *signals*, we suggest following these structural design considerations: *timing*, *frequency* and *tailoring*. First, for a target behavior to take place, the signal needs to occur at the right time. Second, the frequency of the signal will likely affect its efficacy: too frequent prompts may lead to quick saturation and reactance, while too infrequent prompts may be ineffective. While more frequent prompts may be more effective at the initial acquisition of behaviors, less frequent prompting may be more appropriate for behavior maintenance. In addition, prompts that are personalized to a given situation and prompted more often have been proven to be more effective in behavior change than generic and periodical reminders [34].

6 CONCLUSION

In this paper we lay out the design space of technology-mediated nudging through a systematic review of technological prototypes in HCI. Prior frameworks, both in the technological context (e.g. [49], [95]) as well as in marketing

and public policy (e.g., [22][42]) have attempted to review the different cognitive biases and particular implementations of nudges, but did not address the “*how to*” of nudging. This work makes a step forward by linking the why (the cognitive biases) with the how (the exact mechanisms of nudging). Future work will aim at leveraging these insights into a design framework and according tools to support the design of technology-mediated nudges.

ACKNOWLEDGMENTS

This research was partially supported by ARDITI under the scope of the Project M1420-09-5369-FSE-000001, under the scope of the Project Larsys UID/EEA/50009/2019, and by the European Union Co-Inform project (Horizon 2020 Research and Innovation Programme. Grant agreement 770302).

Author contributions AC and EK have conceived and performed the analysis and wrote the paper. AC collected the data. The manuscript has been read by all authors.

REFERENCES

- [1] Alexander T Adams, Jean Costa, Malte F Jung, and Tanzeem Choudhury. 2015. Mindless computing: designing technologies to subtly influence behavior. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 719–730.
- [2] Eytan Adar, Desney S Tan, and Jaime Teevan. 2013. Benevolent deception in human computer interaction. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 1863–1872.
- [3] Elena Agapie, Gene Golovchinsky, and Pernilla Qvarfordt. 2013. Leading people to longer queries. In *Proceedings of the SIGCHI Conference*

- on *Human Factors in Computing Systems*. ACM, 3019–3022.
- [4] Mahdi Nasrullah Al-Ameen, Matthew Wright, and Shannon Scielzo. 2015. Towards Making Random Passwords Memorable: Leveraging Users' Cognitive Ability Through Multiple Cues. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2315–2324.
 - [5] Ernesto Arroyo, Leonardo Bonanni, and Nina Valkanova. 2012. Embedded interaction in a water fountain for motivating behavior change in public space. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 685–688.
 - [6] John A Bargh, Peter M Gollwitzer, Annette Lee-Chai, Kimberly Barn-dollar, and Roman Trötschel. 2001. The automated will: nonconscious activation and pursuit of behavioral goals. *Journal of personality and social psychology* 81, 6 (2001), 1014.
 - [7] Oswald Barral, Gabor Aranyi, Sid Kouider, Alan Lindsay, Hielke Prins, Imtiaz Ahmed, Giulio Jacucci, Paolo Negri, Luciano Gamberini, David Pizzi, et al. 2014. Covert persuasive technologies: bringing subliminal cues to human-computer interaction. In *International Conference on Persuasive Technology*. Springer, 1–12.
 - [8] Ian J Bateman, Alistair Munro, and Gregory L Poe. 2008. Decoy effects in choice experiments and contingent valuation: asymmetric dominance. *Land Economics* 84, 1 (2008), 115–127.
 - [9] Roy F Baumeister, Todd F Heatherton, and Dianne M Tice. 1994. *Losing control: How and why people fail at self-regulation*. Academic press.
 - [10] Henry K Beecher. 1955. The powerful placebo. *Journal of the American Medical Association* 159, 17 (1955), 1602–1606.
 - [11] Robin N Brewer and Jasmine Jones. 2015. Pinterese: exploring reminiscence as an incentive to digital reciprocity for older adults. In *Proceedings of the 18th ACM conference companion on computer supported cooperative work & social computing*. ACM, 243–246.
 - [12] Shun Cai and Yunjie Xu. 2008. Designing product lists for e-commerce: The effects of sorting on consumer decision making. *Intl. Journal of Human-Computer Interaction* 24, 7 (2008), 700–721.
 - [13] Ana Caraban, Maria José Ferreira, Rúben Gouveia, and Evangelos Karapanos. 2015. Social toothbrush: fostering family nudging around tooth brushing habits. In *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers*. ACM, 649–653.
 - [14] Ana Caraban, Evangelos Karapanos, Vítor Teixeira, Sean A Munson, and Pedro Campos. 2017. On the design of subtly: Instilling behavior change during web surfing through subliminal priming. In *International Conference on Persuasive Technology*. Springer, 163–174.
 - [15] Justin Cheng, Chinmay Kulkarni, and Scott Klemmer. 2013. Tools for predicting drop-off in large online classes. In *Proceedings of the 2013 conference on Computer supported cooperative work companion*. ACM, 121–124.
 - [16] Luca Chittaro. 2016. Tailoring Web Pages for Persuasion on Prevention Topics: Message Framing, Color Priming, and Gender. In *International Conference on Persuasive Technology*. Springer, 3–14.
 - [17] Robert B Cialdini. 1987. *Influence*. Vol. 3. A. Michel Port Harcourt.
 - [18] Andy Cockburn, Philip Quinn, and Carl Gutwin. 2015. Examining the peak-end effects of subjective experience. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. ACM, 357–366.
 - [19] Lucas Colusso, Gary Hsieh, and Sean A Munson. 2016. Designing Closeness to Increase Gamers' Performance. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 3020–3024.
 - [20] Matthew Crowley, Aurélie Heitz, Annika Matta, Kevin Mori, and Banny Banerjee. 2011. Behavioral science-informed technology interventions for change in residential energy consumption. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*. ACM, 2209–2214.
 - [21] Tawanna Dillahunt, Olga Lyra, Mary L Barreto, and Evangelos Karapanos. 2017. Reducing children's psychological distance from climate change via eco-feedback technologies. *International Journal of Child-Computer Interaction* 13 (2017), 19–28.
 - [22] Paul Dolan, Michael Hallsworth, David Halpern, Dominic King, Robert Metcalfe, and Ivo Vlaev. 2012. Influencing behaviour: The mindspace way. *Journal of Economic Psychology* 33, 1 (2012), 264–277.
 - [23] Luis Duarte and Luis Carriço. 2013. The cake can be a lie: placebos as persuasive videogame elements. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. ACM, 1113–1118.
 - [24] Dean Eckles, Doug Wightman, Claire Carlson, Attapol Thamrongratanarit, Marcello Bastea-Forte, and BJ Fogg. 2009. Social responses in mobile messaging: influence strategies, self-disclosure, and source orientation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1651–1654.
 - [25] Johan Egebark and Mathias Ekström. 2016. Can indifference make the world greener? *Journal of Environmental Economics and Management* 76 (2016), 1–13.
 - [26] DA Epstein, JH Kang, LR Pina, J. Fogarty, and S. A. Munson. 2016. Reconsidering the device in the drawer: lapses as a design opportunity in personal informatics. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 829–840. <https://dl.acm.org/citation.cfm?id=2971656>
 - [27] Thomas Erickson and Wendy A Kellogg. 2000. Social translucence: an approach to designing systems that support social processes. *ACM transactions on computer-human interaction (TOCHI)* 7, 1 (2000), 59–83.
 - [28] Barbara Fasolo, Raffaella Misuraca, Gary H McClelland, and Maurizio Cardaci. 2006. Animation attracts: The attraction effect in an on-line shopping environment. *Psychology & Marketing* 23, 10 (2006), 799–811.
 - [29] Maria José Ferreira, Ana Karina Caraban, and Evangelos Karapanos. 2014. Breakout: predicting and breaking sedentary behaviour at work. In *CHI'14 Extended Abstracts on Human Factors in Computing Systems*. ACM, 2407–2412.
 - [30] Leon Festinger. 1954. A theory of social comparison processes. *Human relations* 7, 2 (1954), 117–140.
 - [31] Brian J Fogg. 2009. A behavior model for persuasive design. In *Proceedings of the 4th international Conference on Persuasive Technology*. ACM, 40.
 - [32] Alain Forget, Sonia Chiasson, Paul C van Oorschot, and Robert Biddle. 2008. Improving text passwords through persuasion. In *Proceedings of the 4th symposium on Usable privacy and security*. ACM, 1–12.
 - [33] Suzanna E Forwood, Amy L Ahern, Theresa M Marteau, and Susan A Jebb. 2015. Offering within-category food swaps to reduce energy density of food purchases: a study using an experimental online supermarket. *International Journal of Behavioral Nutrition and Physical Activity* 12, 1 (2015), 85.
 - [34] Jillian P Fry and Roni A Neff. 2009. Periodic prompts and reminders in health promotion and health behavior interventions: systematic review. *Journal of medical Internet research* 11, 2 (2009).
 - [35] Luciano Gamberini, Giovanni Petrucci, Andrea Spoto, and Anna Spagnolli. 2007. Embedded persuasive strategies to obtain visitors' data: Comparing reward and reciprocity in an amateur, knowledge-based website. In *International Conference on Persuasive Technology*. Springer, 187–198.
 - [36] Thomas Gilovich, Victoria Husted Medvec, and Kenneth Savitsky. 2000. The spotlight effect in social judgment: An egocentric bias in estimates of the salience of one's own actions and appearance.

- Journal of personality and social psychology* 78, 2 (2000), 211.
- [37] Rúben Gouveia, Evangelos Karapanos, and Marc Hassenzahl. 2015. How do we engage with activity trackers?: a longitudinal study of Habito. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 1305–1316.
 - [38] Rúben Gouveia, Fábio Pereira, Evangelos Karapanos, Sean A Munson, and Marc Hassenzahl. 2016. Exploring the design space of glanceable feedback for physical activity trackers. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 144–155.
 - [39] Rebecca Gulotta, Jodi Forlizzi, Rayoung Yang, and Mark Wah Newman. 2016. Fostering engagement with personal informatics systems. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. ACM, 286–300.
 - [40] Junius Gunaratne and Oded Nov. 2015. Informing and improving retirement saving performance using behavioral economics theory-driven user interfaces. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. ACM, 917–920.
 - [41] Carl Gutwin, Christianne Rooke, Andy Cockburn, Regan L Mandryk, and Benjamin Lafreniere. 2016. Peak-end effects on player experience in casual games. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. ACM, 5608–5619.
 - [42] Pelle Guldberg Hansen and Andreas Maaløe Jespersen. 2013. Nudge and the manipulation of choice: A framework for the responsible use of the nudge approach to behaviour change in public policy. *European Journal of Risk Regulation* 4, 1 (2013), 3–28.
 - [43] Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. 2014. Using personal examples to improve risk communication for security & privacy decisions. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 2647–2656.
 - [44] Marc Hassenzahl and Matthias Laschke. 2015. Pleasurable troublemakers. *The gameful world: Approaches, issues, applications* (2015), 167–195.
 - [45] Sen H Hirano, Robert G Farrell, Catalina M Danis, and Wendy A Kellogg. 2013. WalkMinder: encouraging an active lifestyle using mobile phone interruptions. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. ACM, 1431–1436.
 - [46] Nic Hollinworth, Faustina Hwang, and David T Field. 2013. Using Delboeuf's illusion to improve point and click performance for older adults. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. ACM, 1329–1334.
 - [47] Eleonora Ibragimova, Nick Mueller, Arnold Vermeeren, and Peter Vink. 2015. The smart steering wheel cover: Motivating safe and efficient driving. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 169–169.
 - [48] Nassim Jafarinaimi, Jodi Forlizzi, Amy Hurst, and John Zimmerman. 2005. Breakaway: an ambient display designed to change human behavior. In *CHI'05 extended abstracts on Human factors in computing systems*. ACM, 1945–1948.
 - [49] Eric J Johnson, Suzanne B Shu, Benedict GC Dellaert, Craig Fox, Daniel G Goldstein, Gerald Häubl, Richard P Larrick, John W Payne, Ellen Peters, David Schkade, et al. 2012. Beyond nudges: Tools of a choice architecture. *Marketing Letters* 23, 2 (2012), 487–504.
 - [50] Daniel Kahneman and Patrick Egan. 2011. *Thinking, fast and slow*. Vol. 1. Farrar, Straus and Giroux New York.
 - [51] Daniel Kahneman, Jack L Knetsch, and Richard H Thaler. 1991. Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic perspectives* 5, 1 (1991), 193–206.
 - [52] Yvonne Kammerer and Peter Gerjets. 2014. The role of search result position and source trustworthiness in the selection of web search results when using a list or a grid interface. *International Journal of Human-Computer Interaction* 30, 3 (2014), 177–191.
 - [53] Shipi Kankane, Carlina DiRusso, and Christen Buckley. 2018. Can We Nudge Users Toward Better Password Management?: An Initial Study. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, LBW593.
 - [54] Ted J Kaptchuk, John M Kelley, Lisa A Conboy, Roger B Davis, Catherine E Kerr, Eric E Jacobson, Irving Kirsch, Rosa N Schyner, Bong Hyun Nam, Long T Nguyen, et al. 2008. Components of placebo effect: randomised controlled trial in patients with irritable bowel syndrome. *Bmj* 336, 7651 (2008), 999–1003.
 - [55] Maurits Kaptein, Panos Markopoulos, Boris De Ruyter, and Emile Aarts. 2015. Personalizing persuasive technologies: Explicit and implicit personalization using persuasion profiles. *International Journal of Human-Computer Studies* 77 (2015), 38–51.
 - [56] Evangelos Karapanos. 2015. Sustaining user engagement with behavior-change tools. *interactions* 22, 4 (2015), 48–52.
 - [57] Evangelos Karapanos, Jean-Bernard Martens, and Marc Hassenzahl. 2010. On the retrospective assessment of users' experiences over time: memory or actuality?. In *CHI'10 Extended Abstracts on Human Factors in Computing Systems*. ACM, 4075–4080.
 - [58] Jaejeung Kim, Joonyoung Park, and Uichin Lee. 2016. EcoMeal: a smart tray for promoting healthy dietary habits. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2165–2170.
 - [59] Matthias Laschke, Sarah Diefenbach, Thies Schneider, and Marc Hassenzahl. 2014. Keymoment: initiating behavior change through friendly friction. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational*. ACM, 853–858.
 - [60] Matthias Laschke, Marc Hassenzahl, and Sarah Diefenbach. 2011. Things with attitude: Transformational products. In *Create11 conference*. 1–2.
 - [61] A Lazar, C Koehler, J. Tanenbaum, and D.H. Nguyen. 2015. Why we use and abandon smart devices. *International Joint Conference on Pervasive and Ubiquitous Computing* (2015), 635–646. <https://dl.acm.org/citation.cfm?id=2804288>
 - [62] Dan Ledger and Daniel McCaffrey. 2014. Inside wearables: How the science of human behavior change offers the secret to long-term engagement. *Endeavour Partners* 200, 93 (2014), 1.
 - [63] Min Kyung Lee, Sara Kiesler, and Jodi Forlizzi. 2011. Mining behavioral economics to design persuasive technology for healthy choices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 325–334.
 - [64] Sang-Su Lee, Youn-kyung Lim, and Kun-pyo Lee. 2011. A long-term study of user experience towards interaction designs that support behavior change. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*. ACM, 2065–2070.
 - [65] Birthe A Lehmann, Gretchen B Chapman, Frits ME Franssen, Gerjo Kok, and Robert AC Ruiter. 2016. Changing the default to promote influenza vaccination among health care workers. *Vaccine* 34, 11 (2016), 1389–1392.
 - [66] Q Vera Liao, Wai-Tat Fu, and Sri Shilpa Mamidi. 2015. It is all about perspective: An exploration of mitigating selective exposure with aspect indicators. In *Proceedings of the 33rd annual ACM conference on Human factors in computing systems*. ACM, 1439–1448.
 - [67] Lars Lidén. 2003. Artificial stupidity: The art of intentional mistakes. *AI game programming wisdom* 2 (2003), 41–48.
 - [68] Marilyn McGee-Lennon, Maria Wolters, Ross McLachlan, Stephen Brewster, and Cordelia Hall. 2011. Name that tune: musicons as reminders in the home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2803–2806.
 - [69] Tehila Minkus, Kelvin Liu, and Keith W Ross. 2015. Children seen but not heard: When parents compromise children's online privacy. In

- Proceedings of the 24th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 776–786.
- [70] Andrew Vande Moere. 2007. Towards designing persuasive ambient visualization. In *Issues in the Design & Evaluation of Ambient Information Systems Workshop*. 48–52.
- [71] David Moher, Alessandro Liberati, Jennifer Tetzlaff, and Douglas G Altman. 2009. Preferred reporting items for systematic reviews and meta-analyses: the PRISMA statement. *Annals of internal medicine* 151, 4 (2009), 264–269.
- [72] Sean A Munson, Erin Krupka, Caroline Richardson, and Paul Resnick. 2015. Effects of public commitments and accountability in a technology-supported physical activity intervention. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1135–1144.
- [73] Raymond S Nickerson. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology* 2, 2 (1998), 175.
- [74] Donald A Norman. 2009. THE WAY I SEE IT Memory is more important than actuality. *Interactions* 16, 2 (2009), 24–26.
- [75] Souneil Park, Seungwoo Kang, Sangyoung Chung, and June-hwa Song. 2009. NewsCube: delivering multiple aspects of news to mitigate media bias. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 443–452.
- [76] Charlie Pinder, Jo Vermeulen, Russell Beale, and Robert Hendley. 2015. Subliminal priming of nonconscious goals on smartphones. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*. ACM, 825–830.
- [77] Teppo Räisänen, Harri Oinas-Kukkonen, and Seppo Pahlila. 2008. Finding kairos in quitting smoking: Smokers’ perceptions of warning pictures. In *International Conference on Persuasive Technology*. Springer, 254–257.
- [78] Ilana Ritov and Jonathan Baron. 1992. Status-quo and omission biases. *Journal of risk and uncertainty* 5, 1 (1992), 49–61.
- [79] Yvonne Rogers, William R Hazlewood, Paul Marshall, Nick Dalton, and Susanna Hertrich. 2010. Ambient influence: Can twinkly lights lure and abstract representations trigger behavioral change?. In *Proceedings of the 12th ACM international conference on Ubiquitous computing*. ACM, 261–270.
- [80] William Samuelson and Richard Zeckhauser. 1988. Status quo bias in decision making. *Journal of risk and uncertainty* 1, 1 (1988), 7–59.
- [81] Anuj K Shah and Daniel M Oppenheimer. 2008. Heuristics made easy: An effort-reduction framework. *Psychological bulletin* 134, 2 (2008), 207.
- [82] Patrick C Shih, Kyungsik Han, Erika Shehan Poole, Mary Beth Rosson, and John M Carroll. 2015. Use and adoption challenges of wearable activity trackers. *ICConference 2015 Proceedings* (2015).
- [83] Paul Slovic, Melissa L Finucane, Ellen Peters, and Donald G MacGregor. 2007. The affect heuristic. *European journal of operational research* 177, 3 (2007), 1333–1352.
- [84] Barry M Staw. 1981. The escalation of commitment to a course of action. *Academy of management Review* 6, 4 (1981), 577–587.
- [85] Fritz Strack and Roland Deutsch. 2004. Reflective and impulsive determinants of social behavior. *Personality and social psychology review* 8, 3 (2004), 220–247.
- [86] Erin J Strahan, Steven J Spencer, and Mark P Zanna. 2002. Subliminal priming and persuasion: Striking while the iron is hot. *Journal of Experimental Social Psychology* 38, 6 (2002), 556–568.
- [87] Cass Sunstein. 2013. Impersonal default rules vs. active choices vs. personalized default rules: A triptych. (2013).
- [88] Cass Sunstein and Richard Thaler. 2008. Nudge: Improving decisions about health, wealth, and happiness.
- [89] Cass R Sunstein. 2017. Nudges that fail. *Behavioural Public Policy* 1, 1 (2017), 4–25.
- [90] Anja Thieme, Rob Comber, Julia Miebach, Jack Weeden, Nicole Kraemer, Shaun Lawson, and Patrick Olivier. 2012. We’ve bin watching you: designing for reflection and social persuasion to promote sustainable lifestyles. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2337–2346.
- [91] Yaacov Trope and Nira Liberman. 2010. Construal-level theory of psychological distance. *Psychological review* 117, 2 (2010), 440.
- [92] James Turland, Lynne Coventry, Debora Jeske, Pam Briggs, and Aad van Moorsel. 2015. Nudging towards security: Developing an application for wireless network selection for android phones. In *Proceedings of the 2015 British HCI conference*. ACM, 193–201.
- [93] Amos Tversky and Daniel Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *science* 185, 4157 (1974), 1124–1131.
- [94] Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. 2014. A field trial of privacy nudges for facebook. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 2367–2376.
- [95] Markus Weinmann, Christoph Schneider, and Jan vom Brocke. 2016. Digital nudging. *Business & Information Systems Engineering* 58, 6 (2016), 433–436.
- [96] Neil D Weinstein. 1980. Unrealistic optimism about future life events. *Journal of personality and social psychology* 39, 5 (1980), 806.
- [97] Claes Wohlin. 2014. Guidelines for snowballing in systematic literature studies and a replication in software engineering. In *Proceedings of the 18th international conference on evaluation and assessment in software engineering*. ACM, 38.
- [98] Leslie Wu, Jesse Cirimele, Jonathan Bassen, Kristen Leach, Stuart Card, Larry Chu, Kyle Harrison, and Scott Klemmer. 2013. Head-mounted and multi-surface displays support emergency medical teams. In *Proceedings of the 2013 conference on Computer supported cooperative work companion*. ACM, 279–282.
- [99] Ruud Zaalberg and Cees Midden. 2010. Enhancing human responses to climate change risks through simulated flooding experiences. In *International Conference on Persuasive Technology*. Springer, 205–210.
- [100] Fengyuan Zhu, Ke Fang, and Xiaojuan Ma. 2017. Exploring the Effects of Strategy and Arousal of Cueing in Computer-Human Persuasion. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2276–2283.