Proceedings of the 5th International Conference on Interdisciplinary Humanities and Communication Studies
DOI: 10.54254/2753-7064/2026.39109

To What Extent does Hyper-Nudging Undermine Human Autonomy? A Philosophical Critique of Algorithmic Choice Architecture

Hanqi Haking Liu[1,*], Chengzhe Tang[2]

[1]Tsinglan School, Dongguan, China
[2]Immersion Academy, Irvine, USA
**Corresponding Author. Email: hiboyshello@gmail.com*

Abstract. Personalised recommendation systems have become part of everyday Internet use. Video and shopping platforms use information about what people watch, search for and buy to adjust what appears next. Karen Yeung describes this continuously adaptive digital choice architecture as hypernudging. These systems help users find relevant content and reduce decision-making effort, while their influence on the choice environment raises questions about personal autonomy. This paper examines these questions through Joris Graff's account of the right to one's own reasons, which concerns responsiveness to normative reasons and coherence between actions and broader commitments. External influence does not automatically undermine autonomy. Recommendations can provide useful reasons or introduce interests that users later choose to pursue. Concern arises when personalised influence weakens their ability to assess reasons or repeatedly draws their actions away from commitments they continue to endorse. Kant's distinction between autonomy and heteronomy provides a supporting philosophical perspective.Applying Graff's framework, the paper distinguishes algorithmic assistance from algorithmic interference. Assistance helps users pursue goals they can assess and endorse. Interference becomes troubling when adaptive influence disrupts their ability to act consistently with their considered priorities. The freedom to reject recommendations remains important, though it cannot establish that the wider process respects autonomy. The analysis identifies possible warning signs and design responses without treating them as proof that any particular platform controls its users.

Keywords: hypernudging, personal autonomy, algorithmic choice architecture, right to one's own reasons, personalised recommendation systems

1. Introduction

Personalised recommendation systems influence how people encounter information and opportunities online. A shopping platform arranges products around previous searches. A video application selects material partly on the basis of earlier viewing. The user experiences a manageable stream of options drawn from a much larger collection. This filtering makes services

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2753-7064/2026.39109

easier to use and affects the information available when a decision is made. Research on recommender systems has long treated the reduction of information overload as one of their practical benefits [1].

The choice environment on these platforms is not fixed. It can change as the system receives further information about the user. A person might watch a video out of curiosity and then find similar material appearing more often. Further viewing gives the system additional signals that may affect later recommendations. Yeung's concept of hypernudging captures the ethical significance of adaptive digital environments that guide attention and decisions through continuously updated information [2].

Shopping provides another familiar example. Someone looking for a laptop may compare several models and spend more time on those within a particular price range. The platform can use those interactions to adjust its recommendations. This can help the person find a suitable product without examining every item in the catalogue. The same process may give some products repeated prominence while leaving alternatives harder to discover. A useful service and a powerful means of influence can operate through the same interface.

Users retain the ability to refuse what appears. They can skip a video, close a page or decide against a purchase. This matters because recommendations normally guide behaviour without directly compelling it. The harder question concerns the conditions preceding the final decision. A person may be able to reject an option while having limited awareness of how that option became prominent or why it now seems especially attractive.

Having available choices and directing one's actions are related ideas that need separate examination. A person who clicks voluntarily is still acting within an environment arranged by others. This arrangement may support their purposes or make those purposes harder to follow. The existence of influence alone cannot decide the issue. Family members, teachers and advertising also affect what people notice and want. An account of autonomy must explain which features of influence make a difference.

Graff's work provides the main framework for that explanation. His right to one's own reasons directs attention to how people act and how their reasons fit into their lives. The account does not require complete independence from other people or technology. It asks whether people can respond to considerations that count in favour of an action and whether their decisions remain connected to commitments and projects they recognise as theirs [3].

This distinction improves the analysis of recommendation systems. An algorithm can introduce an unfamiliar book or hobby without taking control away from the user. The person may consider the suggestion and decide that it fits their interests. A system that repeatedly encourages purchases against a user's continuing commitment to save money presents a different concern. The relevant issue is the user's ability to assess and direct that change rather than the fact that the interest first appeared through a recommendation.

Kant supplies a supporting perspective on self-direction. His account connects autonomy to the will's rational legislation of moral principles [4]. Everyday browsing does not usually involve moral self-legislation in that strict sense. This paper uses Kant to explain why the grounds of action matter while relying on Graff for the more direct assessment of ordinary online decisions.

The central research question is:

To what extent can hypernudging undermine Joris Graff's right to one's own reasons and personal autonomy in online environments?

Two further questions guide the discussion:

DOI: 10.54254/2753-7064/2026.39109

1. Where should the normative boundary lie between helpful algorithmic assistance and behavioural influence that threatens the conditions of autonomy identified by Graff?

2. Can the freedom to reject recommendations adequately protect those conditions within continuously adaptive recommendation systems?

The paper argues that a change in behaviour is insufficient evidence of an autonomy violation. Recommendations may help people achieve existing goals or revise them for good reasons. A stronger concern emerges when repeated personalisation weakens their ability to evaluate relevant considerations or persistently disconnects their actions from commitments they have not chosen to revise. Influence on motivation deserves scrutiny when it contributes to either problem.

The distinction between assistance and interference organises this argument. It is a distinction developed here using Graff's account rather than a claim that his article supplies a complete classification of algorithms. Assistance includes recommending laptops that meet a student's needs and helping the student compare them. Interference may occur when the system repeatedly directs attention towards status appeals in ways that make the student's budget or educational purpose increasingly difficult to keep in view.

The same platform may perform both roles during one shopping session. Classification depends on how its influence relates to the person's reasons. Labels such as personalisation or engagement do not settle the ethical judgement.

This paper takes a philosophical and normative approach. Its examples illustrate possible mechanisms without claiming access to a platform's internal model. Technical research and ethical scholarship establish grounds for examining these mechanisms rather than proving that every exposed user loses autonomy [5, 6].

The discussion examines hypernudging and autonomy before considering effects on attention and motivation. It addresses the freedom-to-refuse objection and develops a boundary between assistance and interference. The final discussion considers how platform design could support more considered decisions.

2. Hypernudging and personal autonomy

2.1. From nudging to adaptive choice architecture

Nudging concerns the influence of choice presentation on behaviour while leaving alternatives available. Thaler and Sunstein discuss arrangements that make an option easier to select or more noticeable without prohibiting other choices [7]. Placing healthier food where it is readily visible provides a familiar example. The arrangement affects behaviour through the setting in which a person decides rather than through a direct instruction.

Digital environments extend this possibility through presentation that responds to individual behaviour. A personalised feed changes as information accumulates without waiting for a clear statement of the user's goal. An observed interaction may already become part of its prediction.

Yeung uses hypernudging to describe this dynamic form of influence. Her analysis draws attention to the networked and continually updated character of data-driven choice architecture [2]. The philosophical relevance lies in its ability to adapt the informational setting over time.

Consider a user who watches several videos on the same subject. A platform may use that behaviour to recommend further material on the topic. More viewing can generate further recommendations and create a feedback loop. This example does not require a claim about a particular platform's internal weights. It shows how the options presented and the behaviour used to update them can affect one another.

DOI: 10.54254/2753-7064/2026.39109

A shopping system can operate similarly by using browsing or purchasing information to predict which items a person might consider. Such predictions can save the user effort. A recommendation need not be ethically troubling because it relies on behaviour that the person did not explicitly describe as a preference. Implicit information can help someone find what they want. Its significance depends on how it is interpreted and what the resulting recommendations encourage.

Hypernudging deserves attention because repeated adaptation can change the setting of future decisions. A small interaction today may affect which options become visible tomorrow. This creates opportunities to assist users and opportunities to steer them towards objectives set by the platform. The ethical question concerns the relationship between those objectives and the person's capacity to direct their own activity.

2.2. Attention and the conditions of choice

Recommendation systems select which items receive prominence within collections too large for any person to examine fully. Selection is unavoidable in this setting. The arrangement affects what enters the user's attention and what remains in the background. A person may technically have access to thousands of products while seriously considering only those on the first few screens.

Attention and reasons are different. Seeing an expensive watch does not establish a reason to buy it. Watching a political video does not show agreement with its message. Repeated presentation can make an option familiar or keep it in mind without determining the user's judgement. The argument must explain how exposure interacts with evaluation.

The concern becomes clearer when repeated exposure competes with priorities the user already recognises. Someone intending to spend a short break watching videos may keep encountering content that draws them into another hour of viewing. The activity may be enjoyable and still conflict with their plan to complete other work. A single delay is weak evidence of impaired autonomy. A persistent pattern that the person struggles to assess or change merits closer attention.

Ethical scholarship on recommender systems examines concerns extending beyond predictive accuracy, including autonomy and the way user profiles relate to personal identity [5]. This broader perspective helps explain why a system can predict clicks well without necessarily supporting considered choice. Prediction records what a person is likely to do under given conditions. It does not by itself establish what they regard as worthwhile.

One person may welcome repeated recommendations about a new hobby and another may find them distracting. Assessing autonomy requires attention to users' different commitments and capacities rather than a judgement based solely on content type or frequency.

2.3. Graff and the right to one's own reasons

A narrow account of autonomy might focus on whether a person can act without physical force. Graff's approach asks more of the decision-making process. His first condition concerns responsiveness to normative reasons. A normative reason is a consideration that counts in favour of an action. A motivating reason explains the consideration for which someone actually acts. A factor that causes behaviour need not be a good reason for that behaviour [3].

This distinction matters online. A user may choose a laptop because its features suit their studies. They may also buy a product because repeated status appeals have made it difficult to ignore. The second example requires further examination before judgement. Wanting a product can be reasonable and personal enjoyment can count in its favour. The concern is whether the person can assess those considerations rather than acting through pressures they cannot adequately evaluate.

DOI: 10.54254/2753-7064/2026.39109

Graff's second condition concerns coherence between decisions and broader commitments. An autonomous life need not revolve around one fixed purpose. People can have several interests and change their plans. Their decisions should still bear an intelligible relationship to those commitments or to a revision they can make their own. An appealing immediate option may have reasons in its favour while repeatedly distracting someone from a project they continue to value [3].

The language of ownership used in this paper refers to this practical relationship. It does not mean that people must invent reasons independently or remember every influence on them. A friend's suggestion can become someone's own reason through consideration and acceptance. Dworkin's discussion of autonomy also gives critical reflection an important role in how people relate to their preferences [8]. Reflection helps explain self-direction without requiring a detailed review before every ordinary action.

Graff's approach also avoids making harmful intention the sole basis of criticism. A recommendation process may create risks to autonomy even when no identifiable individual intends to manipulate a particular user [3]. This matters for systems whose influence emerges through optimisation and repeated interactions. The assessment can examine how the process affects users without claiming that the algorithm possesses intentions of its own.

2.4. Assistance and interference

Algorithmic assistance occurs when a recommendation helps someone pursue or reasonably reconsider a purpose they can endorse. The laptop example illustrates the first possibility. Recommending an unfamiliar educational course can illustrate the second. A new interest may support autonomy when the person can assess it and integrate it into their priorities. Existing preferences do not have an automatic claim to remain unchanged.

Algorithmic interference becomes concerning when the system's influence weakens assessment or pulls conduct persistently away from commitments the person still endorses. Repeatedly presenting luxury goods to someone attempting to control unnecessary spending may create this risk. The judgement depends on the pattern of influence and its effects rather than the expense of the goods or the novelty of the desire.

Susser, Roessler and Nissenbaum describe online manipulation in terms of concealed influence that targets decision-making vulnerabilities [9]. Their account supports attention to how influence operates. This paper follows Graff in treating autonomy risks as a broader concern that need not always depend on establishing manipulative intent. The assistance-interference distinction is meant to clarify that concern rather than replace either author's account.

2.5. Kant as a supporting perspective

Kant connects autonomy to rational moral self-legislation. Heteronomy arises when the will takes its determining principle from an object of desire rather than from its own rational legislation [4]. The distinction gives philosophical weight to the grounds on which people act. Merely having several options available cannot capture this conception of autonomy.

The application to recommendation systems remains limited. Enjoying a video or wanting a new laptop does not by itself establish Kantian heteronomy in a way that settles the present argument. Kant's moral account concerns a different level of analysis from Graff's account of personal autonomy. The comparison helps articulate a concern about self-direction while leaving the assessment of online practices to the more suitable framework.

DOI: 10.54254/2753-7064/2026.39109

Yeung explains the adaptive environment in which influence occurs. Graff supplies the central conditions used to evaluate its effects. Kant clarifies the wider philosophical importance of asking what guides action. Keeping these roles distinct allows the analysis to examine ordinary decisions without treating every externally influenced desire as a failure of autonomy.

3. Hypernudging and the reasons that guide action

3.1. Hypernudging reshapes the choice context

Hypernudging usually leaves users able to ignore recommendations. Its influence operates partly through the material available before deliberation. A platform can give some options repeated prominence without issuing commands or removing alternatives. The absence of direct compulsion makes the effect less obvious and creates the need to examine the wider choice environment.

A user might finish an emotionally engaging video and then receive similar content. Further interaction may reinforce that pattern. This is a plausible illustration of a feedback process rather than a description of a verified platform mechanism. It would be unwarranted to assume that replaying a clip always changes a particular internal weight or that every system interprets viewing in the same way.

Chaney, Stewart and Engelhardt use simulations to examine feedback created when recommendation systems learn from behaviour already affected by earlier recommendations. Their results show how this process can produce more homogeneous behaviour without a corresponding improvement in user utility [6]. The study supports concern about feedback effects. It does not demonstrate the loss of personal autonomy in actual users or establish that all recommendation systems produce the same outcome.

The philosophical point is narrower. Behaviour observed within an arranged environment cannot simply be treated as an independent statement of what the user values. A click may reflect genuine interest, momentary curiosity or the prominence of the option. Treating it as a settled preference can help reproduce a pattern the person has never considered adopting.

3.2. From attention to motivation

Attention influences which considerations are available for assessment. Human beings cannot examine every possible option and often work with what is readily visible. Repeated exposure can keep a product or lifestyle near the centre of that assessment. The link to autonomy depends on what follows from this attention and whether the person remains able to evaluate it.

Imagine a shopping system that repeatedly presents status-oriented products after a user spends time viewing them. The user has not expressed an intention to buy luxury goods and is trying to save money. Continued presentation may encourage interest that competes with that goal. This hypothetical example identifies a possible conflict between immediate attraction and an ongoing commitment. It cannot establish the user's vulnerability or the algorithm's intentions without further evidence.

The important question is whether the user can assess the attraction and decide how it fits their priorities. Purchasing a luxury item after considering the expense may be autonomous. Repeatedly buying items against an unrevised savings plan while finding it difficult to step back presents a stronger concern. The origin of a desire provides context and the quality of the person's relation to it carries the normative weight.

DOI: 10.54254/2753-7064/2026.39109

3.3. The boundary between assistance and interference

An algorithm recommending laptops that match a student's stated requirements provides a clear case of assistance. It narrows the search and helps the student compare relevant options. This assessment assumes that the information is sufficiently reliable and that the ranking does not conceal pressures likely to frustrate the student's purpose. Responding to an explicit query is encouraging evidence rather than a guarantee of harmlessness.

A more troubling case involves repeated recommendations that exploit a person's susceptibility to status appeals and displace their continuing priorities. This example fits the concern described in scholarship on covert influence and decision-making vulnerabilities [9]. Establishing that it occurs on a particular platform would require evidence about the system and the user's experience.

The boundary cannot be drawn simply between old and new desires. Learning often changes what people want. A useful recommendation can introduce a goal the person had never considered. Assistance and interference differ in whether users remain able to assess those changes and connect their conduct to commitments they recognise. The system's contribution to motivation matters when it undermines that relationship.

3.4. Self-direction and externally shaped motivation

A desire may feel personal even when repeated digital influence has helped strengthen it. That experience cannot settle whether the process respects autonomy. A person may sincerely welcome an interest developed online or feel unable to pursue a different priority they still value. Both possibilities require room for the person's own assessment.

Kant's concern with the determining grounds of the will offers a limited parallel [4]. The present claim remains about ordinary self-direction rather than strict moral self-legislation. Adaptive recommendation raises an autonomy risk when its influence makes reasoned assessment or coherence with considered commitments harder to sustain. The analysis does not require the stronger claim that algorithms take over human decision-making.

4. Counterargument and rebuttal

4.1. The freedom to reject recommendations

The strongest objection begins with users' continuing capacity to refuse. They can skip posts, leave shopping pages and search elsewhere. Recommendations influence the presentation of options while the person makes the final decision. A defender of personalisation may argue that this influence is compatible with autonomy and often helps people use their time more effectively.

Sunstein argues that appropriate nudges can support agency and consumer freedom [10]. His position gives the objection a serious basis without establishing that every adaptive commercial recommendation is appropriate. A fair reconstruction should defend the benefits of guided choice rather than assume that any voluntary click proves complete autonomy.

The objection has three main elements:

1. Users retain meaningful decision-making authority when refusal and alternative action remain practical possibilities.

2. Every social environment affects attention and desire, so algorithmic influence needs a specific explanation of what makes it ethically different.

DOI: 10.54254/2753-7064/2026.39109

3. Useful personalisation and accessible controls can support users' purposes more effectively than requiring them to manage all available information alone.

These claims challenge any blanket condemnation of recommendations. A system that helps someone locate useful material may strengthen their ability to act on their reasons. The person can also choose to delegate some filtering to a service because examining everything would be impractical. Autonomy should allow this kind of considered reliance.

The objection is especially persuasive where users understand recommendations and can readily change settings. The critic must explain how a particular pattern threatens autonomy rather than relying on the fact that information has been selected.

4.2. Why refusal does not settle the question

The freedom to reject protects an important part of agency. It does not establish that the conditions under which a person forms and follows reasons are adequate. A user can remain free to decline a purchase while encountering a stream that repeatedly draws attention away from a financial commitment. The final click leaves this broader relationship unresolved.

Suppose a shopping catalogue contains ten relevant options and the interface repeatedly foregrounds one through emotionally appealing presentation. The user can scroll further and refuse the prominent item. That possibility matters. It does not tell us whether repeated exposure has made alternative considerations difficult to assess or encouraged conduct disconnected from the user's continuing plans. Those are separate questions about the process of choice.

A skip button cannot answer them by itself. The person may use it easily or may repeatedly struggle to do so despite intending to move on. An assessment needs evidence about the practical pattern and the person's priorities. Occasional distraction should not be inflated into a loss of autonomy. A design that persistently makes considered control difficult deserves more scrutiny.

The comparison with ordinary social influence also needs care. Advertising and personal persuasion can be manipulative or responsive to individual reactions. Digital systems do not introduce influence into an otherwise neutral world. Their relevant features may include scale, continuous adaptation and limited visibility into how responses affect later presentation. These features can intensify familiar problems without making every digital interaction fundamentally different.

Exit and preference controls help when users can find, understand and use them. Their effectiveness cannot be assumed from their technical existence. Nor can this paper assume that every platform hides them. The appropriate question is whether the controls provide practical influence over the recommendation process and whether using them allows the person to pursue their considered purposes.

The objection retains an important insight. Personalisation can save effort and help users follow goals they endorse. The rebuttal shows that refusal alone is an incomplete safeguard. Evaluation must consider both the ability to decline individual recommendations and the user's continuing capacity to assess reasons and direct activity across repeated interactions.

5. Discussion

5.1. Drawing the normative boundary

The analysis supports three categories for examining algorithmic influence. They are practical distinctions proposed in this paper and should not be mistaken for a classification directly

DOI: 10.54254/2753-7064/2026.39109

established by Graff:

1. Algorithmic assistance helps users pursue or thoughtfully revise goals they can endorse. Relevant recommendations and understandable controls support this role.

2. Ordinary algorithmic influence changes what users notice or consider while leaving their capacity to evaluate reasons and maintain considered commitments substantially intact.

3. Autonomy-threatening interference creates a credible risk to reasoned assessment or repeatedly separates conduct from commitments users continue to endorse.

The third category requires more evidence than changed behaviour. A person might spend more time on a hobby because the platform helped them discover something worthwhile. Another might repeatedly spend time against a plan they still wish to follow. The observed duration can be similar and the relationship to self-direction quite different.

Graff's account directs attention to reasons and coherence rather than offering a simple measurement of algorithmic control [3]. The application here asks whether a system's influence makes it harder for users to assess relevant considerations or organise conduct around their own projects. These conditions can be threatened through persistent distraction as well as through the strengthening of new desires.

Possible warning signs include strong reliance on behaviour that is difficult to interpret as considered preference and feedback that repeatedly narrows exposure despite the user's attempts to change it. Another is an objective that rewards continued engagement while giving little practical weight to user-specified limits. These features justify investigation when they occur together or produce patterns users struggle to redirect.

The warning signs are neither necessary nor sufficient conditions for an autonomy violation. Implicit information can be helpful and specialised recommendations can suit a user's purpose. Diverse content can still exert unwelcome pressure. The judgement requires an account of the interaction between system design and the user's ability to pursue or reconsider their commitments.

5.2. Awareness and reflective capacity

Awareness can help users interpret recommendations. Someone who understands that a feed responds to viewing may recognise why a topic keeps returning and reconsider whether to engage with it. This understanding can create space for assessment. It does not guarantee that the person will find repeated distractions easy to manage.

Kozyreva, Lewandowsky and Hertwig discuss cognitive tools intended to strengthen agency and resilience in digital environments [11]. Their approach supports attention to users' abilities as well as to platform design. The implication for this paper is that a brief explanation of personalisation may be useful without providing all the support needed for considered control.

Reflection concerns whether people can assess the role an activity plays in their lives. A user might ask whether repeated shopping fits their budget or whether another hour of videos is how they want to spend the evening. They do not need to perform a philosophical analysis before each click. They need reasonable opportunities to notice a conflict and act on their judgement.

That opportunity can depend on the environment. Controls that are easy to use may let people pause a recommendation pattern or change the basis on which content is selected. Prompts to reconsider continued use may help some users and irritate others. Their value should be assessed through practical experience rather than assumed from a label such as autonomy support.

Responsibility cannot rest entirely on individual vigilance. Users may face complex systems during ordinary activities when attention is already limited. A platform that supplies useful

DOI: 10.54254/2753-7064/2026.39109

personalisation can also make it easier to understand and redirect that service. This creates a shared task involving user competence and design that leaves room for considered choice.

5.3. Limits of the argument

The analysis has several limitations. Human preferences develop through social relationships and cultural exposure, so there is no wholly uninfluenced baseline against which algorithmic effects can be measured. An interest acquired online cannot be judged inauthentic merely because another source helped produce it. The assessment concerns how the person relates to that interest.

The causal role of recommendations is also difficult to isolate. A change in behaviour may reflect earlier preferences, platform influence or other developments in the person's life. Feedback between behaviour and presentation makes the problem harder because each may affect the other. The simulation evidence discussed earlier identifies a possible mechanism without resolving these questions for actual users [6].

Reflection is not an infallible test. People may give incomplete explanations of their actions or revise their commitments gradually. Asking whether a person endorses a choice remains useful, though an immediate answer cannot establish the entire history or quality of the process. Autonomy assessment needs sensitivity to continuing patterns and opportunities for revision.

This paper also offers a normative argument rather than experimental evidence. It identifies conditions that deserve scrutiny and proposes ways to examine them. It cannot assign an effect size to hypernudging or conclude that a named platform violates its users' autonomy. Interviews, behavioural research and access to system information would be needed to develop those claims.

Certain recommendation practices can create wrongful risks even where a completed violation cannot be proved. Further investigation should ask which users face those risks and whether design changes improve their capacity to direct activity. The philosophical account supplies a standard for interpreting the answers.

5.4. Implications for algorithmic choice architecture

Platform evaluation should examine more than whether users retain visible options to skip or leave. It should consider whether personalisation supports the assessment of reasons and the pursuit of considered commitments. This proposal extends the ethical concern developed in the paper. The suggested measures are starting points for evaluation rather than claims about existing legal requirements.

One dimension is understandable information about recommendation inputs. Users could be told whether a suggestion mainly reflects a stated preference, recent activity or some other broad basis. Controls could allow them to reduce the role of particular behavioural signals. The aim is practical understanding rather than requiring users to read the technical details of a model.

A second dimension concerns control over feedback. Users could pause learning from recent activity, reset recommendations or request a broader range of material. These options may help someone change a pattern that no longer suits their interests. Their effectiveness should be evaluated rather than presumed. A fixed diversity quota cannot guarantee autonomous choice and may frustrate a person seeking specialised information.

A third dimension concerns the objectives used to evaluate success. Continued engagement and purchases reveal something about behaviour, though they provide an incomplete account of benefit. Evaluation could also examine whether users can follow self-selected limits and whether

DOI: 10.54254/2753-7064/2026.39109

recommendations support purposes they identify as worthwhile. Gathering such evidence should itself respect privacy and allow users to revise their aims.

An explanation has limited value if the person cannot alter the process. A control has limited value if its effect is unclear. Measures of success can reward unwanted pressure when they ignore users' considered priorities. The relationship between these features deserves empirical study.

The design challenge is to retain useful personalisation while giving users practical influence over how it serves them. A recommendation system can help people explore unfamiliar options and make decisions with less effort. Its benefits are more defensible when users can assess that assistance and redirect it as their purposes change. This is a concrete way to connect the right to one's own reasons with everyday platform use.

6. Conclusion

Hypernudging does not undermine autonomy merely because it changes behaviour. Personalised recommendations can reduce search effort and help people discover worthwhile options. External influence is part of ordinary decision-making and a newly formed interest can become part of an autonomous life. The question concerns how people remain able to assess and direct that influence.

Graff's account provides two central considerations: responsiveness to normative reasons and coherence between conduct and broader commitments. The distinction between assistance and interference developed here applies those considerations to adaptive recommendation. A system threatens autonomy when its influence weakens assessment or persistently draws actions away from commitments the user has not chosen to revise. Its causal contribution to a desire alone does not establish a violation.

The freedom to refuse remains a valuable safeguard. Users who can easily decline recommendations and redirect the service have greater scope to act on their purposes. Formal refusal cannot by itself establish that a continuously adaptive environment protects the conditions of self-direction. The process preceding a click and the pattern across repeated decisions also matter.

Kant offers a supporting reminder that agency involves the grounds of action as well as available options His moral conception is not a direct test for ordinary browsing. Graff supplies the more suitable framework for examining whether online influence remains compatible with a person's own projects and reasons.

The ethical assessment should focus on credible risks and practical opportunities for considered control. Understandable recommendation inputs, usable feedback controls and evaluation linked to users' purposes offer possible responses. These proposals require further testing. They preserve the paper's central claim that effective personalisation should help people act from reasons they can assess and integrate into their lives.

Author contribution

All the authors contributed equally and their names were listed in alphabetical order.

References

[1] Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering, 17*(6), 734–749. https://doi.org/10.1109/TKDE.2005.99

[2] Yeung, K. (2017). 'Hypernudge': Big data as a mode of regulation by design. *Information, Communication & Society, 20*(1), 118–136. https://doi.org/10.1080/1369118X.2016.1186713

DOI: 10.54254/2753-7064/2026.39109

[3] Graff, J. (2026). The right to one's own reasons: Autonomy and online behavioural influence. *Moral Philosophy and Politics, 13*(1), 41–70. https://doi.org/10.1515/mopp-2025-0034
[4] Kant, I. (1998). *Groundwork of the metaphysics of morals* (M. Gregor, Ed. & Trans.). Cambridge University Press. https://doi.org/10.1017/CBO9780511809590 (Original work published 1785)
[5] Milano, S., Taddeo, M., & Floridi, L. (2020). Recommender systems and their ethical challenges. *AI & Society, 35*(4), 957–967. https://doi.org/10.1007/s00146-020-00950-y
[6] Chaney, A. J. B., Stewart, B. M., & Engelhardt, B. E. (2018). How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM Conference on Recommender Systems* (pp. 224–232). Association for Computing Machinery. https://doi.org/10.1145/3240323.3240370
[7] Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.
[8] Dworkin, G. (1988). *The theory and practice of autonomy*. Cambridge University Press. https://doi.org/10.1017/CBO9780511625206
[9] Susser, D., Roessler, B., & Nissenbaum, H. (2019). Technology, autonomy, and manipulation. *Internet Policy Review, 8*(2). https://doi.org/10.14763/2019.2.1410
[10] Sunstein, C. R. (2015). Nudges do not undermine human agency. *Journal of Consumer Policy, 38*(3), 207–210. https://doi.org/10.1007/s10603-015-9289-1
[11] Kozyreva, A., Lewandowsky, S., & Hertwig, R. (2020). Citizens versus the internet: Confronting digital challenges with cognitive tools. *Psychological Science in the Public Interest, 21*(3), 103–156. https://doi.org/10.1177/1529100620946707