Balancing Personalization and Public Values: Legitimacy and Design of Algorithmic News Recommendations in Public Service Media

Laura Finarelli[1,4][0009−0000−8349−420X], Johan Rochel[2][0000−0002−8123−4057], Florian Evéquoz[1][0000−0002−6052−2252], and Gianluca Rizzo[1,3][0000−0001−7129−4972]

[1] HES-SO Valais, Switzerland
[2] University of Fribourg & EPFL, Switzerland
[3] Universita' di Torino, Italy
[4] Technical University of Berlin, Germany
laura.finarelli@hevs.ch

Abstract. This paper addresses the normative challenges of algorithmic recommendation systems in public service media (PSM), proposing a framework that aligns democratic mandates with digital engagement via user needs categorization. Institutions like Radio France and the BBC use algorithms to counter news avoidance and filter bubbles through civically weighted content and discoverability features. Yet, their non-commercial missions create tensions between user autonomy, transparency, and societal value. The study introduces a paradigm shift: modeling user behavior through why-oriented needs (e.g., civic awareness, informational gaps) rather than how-focused engagement metrics. A 2023 case study with Switzerland's RTS develops three nudging scenarios using community detection and needs-based clustering. These inform a tripartite legitimacy framework: 1) alignment with public service duties, 2) ethical user guidance, and 3) systemic risk mitigation via participatory design. Findings show that needs-aware systems require multidimensional profiling balancing explorability and explainability, distinct from commercial logic. Bridging data science and nudging ethics, this work advances interdisciplinary strategies for operationalizing PSM's dual mandate: respecting individual agency while fostering democratic citizenship.

Keywords: Nudging · User Needs · Community Detection · Public Service Media · News Provider · Recommendation System · Ethics · Big Data

1 Introduction

Public service media (PSM) are currently facing the opportunities and tensions accompanying the introduction of algorithmic recommendation systems [23, 7]. These recommendation systems are primarily used to calibrate and optimize the content highlighted through public service applications. For instance, since the end of 2022, Radio France has been using an algorithmic recommendation system to disseminate the content available on its application (shows, podcasts, etc.) to specific audiences to increase "discoverability" [17]. If they consent, app users are invited to discover content beyond their usual preferences. PSM may have non-commercial objectives when they use algorithmic recommendation systems, such as aligning content use to the public mission and values of the PSM operator itself. For instance, the BBC's normative framework emphasizes diversity of exposure over purely popularity-driven recommendations, requiring algorithms to weight content based on its civic importance rather than predicted virality [37]. In line with this perspective, algorithmic recommendation systems are considered promising tools to tackle two major challenges in news consumption: on the one hand, news avoidance [26, 25], and on the other, the effects of filter bubbles. From the perspective of society, filter bubbles may contribute to widespread polarization, as individuals are increasingly exposed to information that reinforces their existing beliefs and biases [1]. Radio France's ambition of "discoverability" may be seen as an attempt to face these two challenges. Giving users content that might surprise them is expected to spark interest (against news avoidance and fatigue) and to open up new perspectives on relevant issues (against filter bubble effects) [19].

This study aims to balance the previously introduced negative trends by introducing a new tool designed to recommend news articles based on the underlying reasons behind users' news consumption, focusing on the "why" rather than the "what." Additionally, it addresses the legitimacy of PSM in utilizing such algorithmic recommendation systems. In 2022, the Committee of Ministers of the Council of Europe urged member states to establish governance frameworks to address the risks associated with algorithmic content curation, selection, and prioritization [8]. However, before such mechanisms can be designed effectively, it is essential to clearly identify the specific ethical and normative challenges PSM face when implementing these systems. In taking up this task, this interdisciplinary contribution considers the use case of algorithmic nudging of news articles based on user needs [35], which enables the modeling of user behavior and preferences not only by documenting *how* but also *why* they consume news, thereby enabling publishers and journalists to make more accurate and personalized recommendations.

User needs were introduced in 2016 by the journalist Dmitry Shishkin (see Figure 1), from the BBC, to sum up the reasons why users read certain news and to categorize the tone of the articles [36]. The main goal behind this approach is to put the reader at the center of content strategy, by telling a story from an angle that the audience values [35].

This innovative approach to algorithmic nudging represents a paradigm shift in recommendation systems, moving beyond traditional engagement-driven algorithms to incorporate elements of behavioral science and cognitive heuristics based on a user needs categorization of user patterns. The shift toward needs-based recommendations introduces several open ethical questions about how to transparently guide choices without manipulation. [28] emphasizes the risks of embedding human biases into algorithms, while [20] advocates for user participation in system design to maintain trust.

However, despite the high potential for more accurate and personalized recommendations, so far, no news recommendation algorithm accounts for user needs. Thus, to tackle the above-mentioned issues, we first describe an innovative approach to user nudging based on modeling content consumption patterns by means of community detection algorithms and user needs categorization. This approach draws upon a study conducted with the French-speaking branch of the Swiss public media company (RTS - Radio Télévision Suisse). The case study relies upon existing news content and real engagement data gathered between March and May 2023 from the platforms of RTS.

Based on this approach, we have defined three nudging scenarios on which to base our analysis of the normative challenges related to the use of behavior-based recommendation systems in the context of a PSM.

Secondly, in order to define algorithmic curation by PSM as digital nudges that balance engagement and societal good, and to justify principles that should guide this use, we have gathered these principles in a three-component legitimacy framework: The objectives of the recommendation system, the implementation approaches, and the analysis of aggregated risks.

Our work thus contributes to the literature on algorithmic recommendation in the field of PSM while discussing the legitimacy of the proposed tool. A significant part of this literature draws upon social sciences methods (e.g., interviews in newsrooms)[39, 33]. This contribution complements this scholarship by drawing upon data science methods and by addressing ethical challenges from the broader perspective of the ethics of nudging [24, 34].

This paper is organized as follows. Section 2 introduces the data model and methods employed to categorize user behavior. This includes an analysis of content consumption patterns. Section 3 presents the proposed recommendation tool, built upon the definition of three nudging scenarios, illustrating how algorithmic recommendations can influence user engagement. Section 4 illustrates the use of the recommendation tool on concrete examples.

Drawing on the three nudging scenarios, Section 5 delves into the normative and ethical challenges posed by algorithmic recommendation systems in the context of PSM, with a specific focus on our proposed approach. It explores a legitimacy framework based on three key components: The system's objectives, the modalities of implementation, and the analysis of aggregated risks. The discussion emphasizes the need for ethical considerations in the design and use of such systems, particularly concerning their impact on news diversity and democratic engagement. Finally, the paper concludes by reflecting on the broader implications of these findings for both public and commercial media, highlighting the role of algorithmic systems in shaping future news consumption patterns and the ethical responsibilities associated with their use.

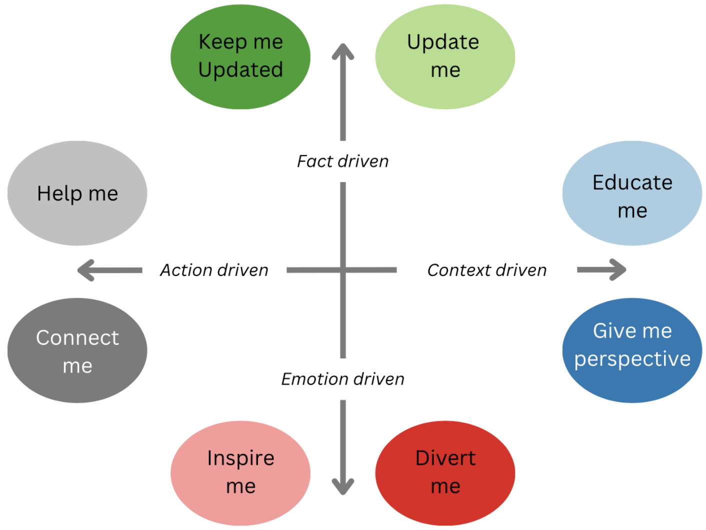

Fig. 1: User needs model 2.0 [36]. The proposed user needs are grouped into four main poles.

2 Datasets and Methodology

This section presents the methodology: first, the datasets (consumption logs, articles metadata); next, the graph-based community detection model and user cluster analysis; and finally, the three nudging scenarios integrated into the recommendation system.

2.1 Datasets

Digital media organizations manage detailed content classifications, including articles and rich metadata. They also collect audience data to track user interactions with published content, supporting their core activities through insights into both the content and its contextual and engagement elements. This study draws on two datasets from RTS: a structured dataset of 3,982 online articles published between March 1st and May 31st, 2023, and server logs recording user engagement during the same period. Such data is commonly available in digital newsrooms. The article dataset includes metadata for each article, notably one or two user needs assigned from six categories: Update me, Divert me, Inspire me, Keep me on trend, Educate me, and Give me perspective[5]. The topics of the articles can be retrieved from the logs dataset and correspond to the main topic sections available on the RTS website. After removing the articles without an associated user need, we obtained 3.713 articles. The statistics related to the user need frequencies per topic, and the topics are summarized in Figure 2. As evident from the plot, the majority of the articles are associated with the user need Update me, followed by Give me perspective and Educate me. These statistics are consistent with the most important user needs highlighted in [25].
The logs dataset consists of server logs containing each access to an URL of RTS website accessed during the period from the 1st of March to the 31st of May, 2023. The logs were filtered to only include rows where an article from the article dataset was present. The distribution of topics based on the visits is illustrated in Figure 3. The plot also displays the relative distribution of articles and their corresponding access frequencies. According to the suggestion of the RTS editors, we confirmed the presence of bots in the logs during the cleaning process. Some of these bots were easily identified by their agent IDs and were filtered out. Additionally, to avoid biases from consistently refreshing the same article, we only accounted for the user-article association once when joining the datasets.
The logs and articles datasets have been joined using the unique ID of the article that is present in both

[5] User need model 2.0 includes two more user needs that were not available in our dataset. However, our proposed framework can be easily extended.

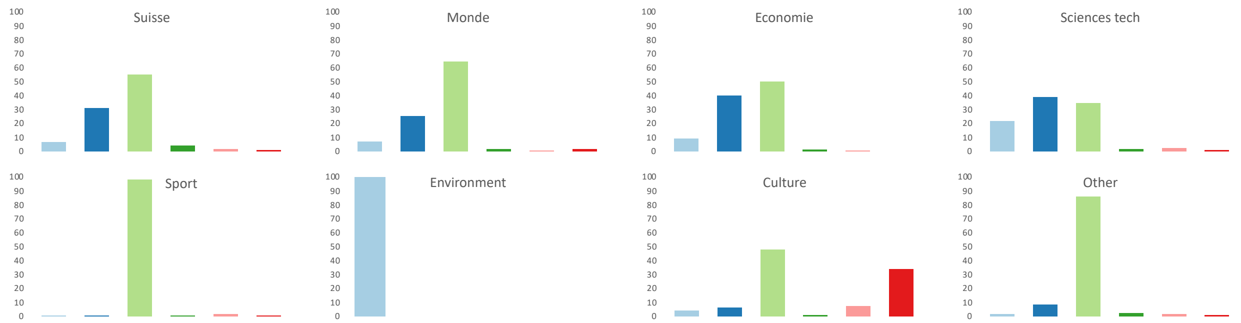

Fig. 2: Proportion of visits by user need, for each topic. Colors match user needs as follows: ■ Educate me; ■ Give me perspective; ■ Update me; ■ Keep me on trend; ■ Inspire me; ■ Divert me.

datasets. Of particular value for our study are the couples composed by the topic and the user need. Taken together, they allow for a better characterization of the articles and users' interests. Only 36 out of 42 possible combinations are present. The topic *économie* is never associated with the user need Divert me, and the topic *environnement* is only present in association with Educate me. The distribution of visits by combinations/topic is visible in Figure 3.

The final joined dataset associates each user with the diverse couples (topic, user need) corresponding to the pages they visited, along with the number of visits associated with each couple.

2.2 Clustering and community detection

To better understand audience news consumption and design an effective recommendation system, we identified user communities based on content consumption patterns. This involved building a user network graph, where nodes represent users and edges reflect behavioral similarity. Graph-based clustering reveals tightly connected communities and captures complex, non-linear relationships in the data, offering insights into user behavior and supporting personalized recommendations [16].

We represent the data as a graph, where each user is a node and is connected to other users. One of the advantages of using graphs is their ability to identify connected components (clusters). This means the organization of vertices in communities, where there are many edges joining vertices within the same cluster and comparatively few edges joining vertices from different clusters [16]. Furthermore, graph techniques can detect non-convex and non-linear relations between data. Our data can be translated to a biadjacency matrix B where the rows are the users and the columns are the combinations. The entries $B(i,k)$, $i = 1, .., \#users$, $k = 1, .., \#combinations$, are the numbers of occurrences of user i and combination k, normalized between 0 and 1 with respect to the total number of articles read by each user. Our study did not account for occasional users, identified as those with less than 30 distinct articles read in the 3-month-long observation window. That is, users with lower values are users who, on average, read no more than 1 article every 3 days and are not considered relevant for our partner and thus our analysis. Furthermore, we did not account for users whose total number of visited articles falls below the minimum amount sufficient to characterize their interest. The final number of considered users is 18925.

In the graph, two users are connected by an edge if they share at least one combination (topic, user need). The weight of the edge is determined by the cosine distance of frequencies for each shared combination, and the frequency is modulated using the term frequency-inverse document frequency (TF-IDF) to give more value to less frequent, and thus more informative combinations. In such a way, users with more combinations in common will be generally closer to each other, i.e., have a higher similarity measure. Furthermore, the adopted measure ensures a faster way to generate the graph of users since the matrix of similarities is symmetric. Given the high number of edges, resulting in a very dense and difficult-to-handle graph, we decided to filter the edges with weights (i.e., similarity) lower than a certain threshold. The threshold, set to 0.8, is chosen as a good compromise between accuracy and computational efficiency. We compared four different community detection techniques: Louvain, FastGreedy, Label Propagation, and Leiden [22, 9, 41]. Modularity and conductance [29] were adopted as metrics to evaluate the quality of the results. Modularity

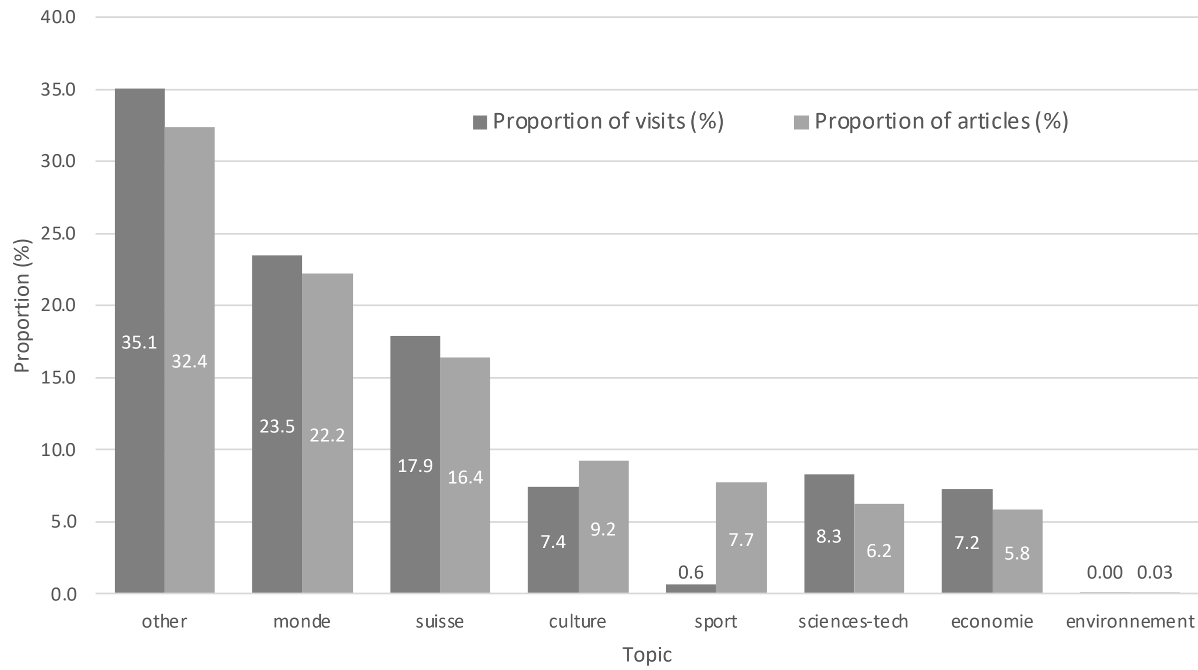

Fig. 3: Distribution of article visits and published articles by topic in the dataset. Except for the topic *sport*, the proportion of visits is coherent with the publishing trend.

measures the strength of the division of a network into communities by comparing the density of edges within communities to the density of edges between communities [42, 16]. Conductance measures the quality of a partition by examining the ratio of edges that are external to the community versus internal.

We found that the Leiden algorithm performed the best in terms of these metrics. The resulting communities and their relationship can be visually seen in Figure 4. The thickness of the edges is proportional to the number of connections going from one community to the connected one. The weight is normalized to the smaller community size. By normalizing by the size of the smaller community, it is possible to account for the fact that smaller communities have fewer potential connections. This helps to ensure that the weight reflects the relative strength of the connection rather than just the absolute number of interactions. For communities 14 and 15, no similarity greater than 0.8 between users in other communities is present so they result *isolated* in the Figure.

To better characterize communities, we retrieved the four most common combinations (topic, user needs) in each community. These are presented in Table 1. To assess the potential of including user needs in the analysis, we compare the results with the community detection that only considers the topics of the articles. The metrics, algorithms, and data preprocessing are the same as in the more general case. As previously experienced, 5 out of 7 topics were common to almost all users, while only *sport* and *environnement* were exclusive to some users. This bias in the data is mainly due to the unbalanced distribution of topics among articles, as it is evident in Figures 3 and 2. Our attempt to find communities was ineffective and of poor quality. The most effective algorithm was the fast greedy algorithm, which identified 3 communities with a modularity of 0.03 and a conductance of 0.64. These values indicate a low community structure and high randomness. As anticipated, considering user needs in our study allowed for a more detailed understanding of users and improved results.

3 Recommendation tool

To prevent filter bubbles and enhance content discovery, we propose recommending content that encourages users to explore across community boundaries. This approach aligns with digital ”nudging” (see Section 5).

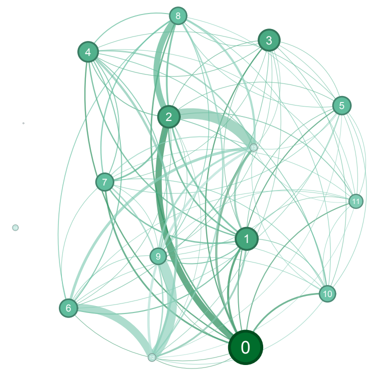

Fig. 4: Communities detected by the Leiden algorithm. Labels on the nodes identify the communities. Node size and color match the number of members of that community: the bigger and darker the node, the bigger the community. The thickness of the edges is proportional to the number of connections going from one community to the connected one. The weight is normalized with respect to the smaller community size. One community has a below-threshold number of links to other communities and appears as not connected to the rest of the graph.

We define three nudging scenarios for public service media: Recommending content under the user need Give me perspective, without a specific theme; recommending Swiss political news aligned with Give me perspective; suggesting content distant from the user's current community cluster. These scenarios guide users toward diverse content journeys and will be referred as:

1. *Give-me-perspective nudging*: The user can choose their favorite topic while assuming Give me perspective as the user need;
2. *Citizen-perfectionist nudging*: The topic is *Suisse* and user need Give me perspective;
3. *Random discovery*: The user is randomly suggested an article frequently read by another community, both to surprise him and let him explore out of his comfort zone.

These three scenarios highlight key issues in algorithmic recommendations. On one hand, they raise ethical questions about nudging (see Section 5); on the other, they propose technical challenges. The remainder of this section outlines their implementation, illustrated through a use case in Section 4.
The general scheme for our tool can be summarized as follows:

– Input: a news consumer/user characterized by an array of features, i.e., the couples (topic, user need), and the values of each feature are the times the combinations occur;
– Community Assignment: the user is assigned to the closest (most similar), community through the pre-trained, feature-based, clustering algorithm;
– Direction to the target Community: The user is directed towards the target community, passing through the intermediate communities, determined using distance metrics in the feature space;
– Output: a list of certain kinds of articles that the intermediate and final communities generally engage with.

The proposed approach is schematized in Figure 5. The primary goal is to gradually guide the user through intermediate communities from their current community to the target community, making the nudging process *smoother*.

Community	1st Feature	2nd Feature	3rd Feature	4th Feature
0	culture Inspire me	suisse Inspire me	monde Divert me	culture Educate me
1	sciences-tech Inspire me	monde Divert me	monde Keep me on trend	culture Inspire me
2	culture Inspire me	culture Educate me	monde Keep me on trend	sciences-tech Inspire me
3	economie Keep me on trend	monde Divert me	monde Keep me on trend	culture Inspire me
4	sciences-tech Keep me on trend	monde Divert me	monde Keep me on trend	culture Inspire me
5	sciences-tech Divert me	monde Divert me	monde Keep me on trend	sciences-tech Inspire me
6	monde Divert me	monde Keep me on trend	culture Inspire me	sciences-tech Inspire me
7	suisse Divert me	monde Divert me	culture Educate me	sciences-tech Inspire me
8	culture Educate me	monde Divert me	monde Keep me on trend	culture Inspire me
9	monde Keep me on trend	monde Divert me	culture Inspire me	culture Educate me
10	monde Inspire me	monde Divert me	monde Keep me on trend	culture Inspire me
11	culture Keep me on trend	monde Divert me	monde Keep me on trend	culture Educate me
12	monde Divert me	monde Keep me on trend	culture Inspire me	sciences-tech Keep me on trend
13	culture Inspire me	monde Divert me	monde Keep me on trend	culture Educate me
14	sport Update me	culture Inspire me	suisse Inspire me	culture Educate me
15	culture Inspire me	economie Inspire me	monde Divert me	culture Educate me

Table 1: Top four features for each community.

Combination	Community label
monde Give me perspective	0
suisse Give me perspective	15
sciences-tech Give me perspective	4
economie Give me perspective	15
culture Give me perspective	0
sport Give me perspective	10

Table 2: All the possible combinations with the user need Give me perspective associated with the community where they are more frequently read.

4 Use Case

This section offers an example of how our proposed nudging tool operates.

When a new user visits the news site, the recommendation tool activates after a set time or threshold (e.g., 30 articles), balancing accuracy and computational cost. Based on user preferences (collected with consent), one of three nudging strategies is applied. We illustrate this with two users: one real and one synthetic (with uniformly distributed feature values). The tool links users to the nearest community based on similarity to the average community profile (i.e., the center of the cluster). It then recommends content by gradually guiding users toward a target community. To identify target communities for the Give me perspective user need, we extracted the most frequent topic-user need combinations per community. These associations are summarized in Table 2. Notably, no articles matched the topic *environment* with the Give me perspective need.

We tested all three nudging strategies and found they yield distinct recommendation patterns.

The first user, randomly extracted from the dataset, belongs to community 4 and is characterized by the combinations *culture* Divert me, *culture* Update me, *economie* Educate me, *economie* Update me, *monde* Give me perspective, *monde* Update me, *sciences-tech* Educate me, *sciences-tech* Update me, *suisse* Give me perspective, *suisse* Educate me, *suisse* Keep me on trend, and *suisse* Update me with frequencies $1, 1, 1, 1, 4, 6, 1, 4, 6, 1, 1$ and 7 respectively. The proposed tool correctly associates the user with community 4.

According to the value of *nudging* (default 1) and *favouriteTopic* (default *None*), users are suggested certain types of articles to guide them to the target community, passing through intermediate communities:

1. "Give-me-perspective nudging" ($nudging = 1$, $favouriteTopic = sport$): the user should be lead to community 10. The suggested combinations (corresponding to articles) proposed by our tool to reach the

Legal framework	**Specific objectives**	**Relevant user needs**	**Possible architectural measures**
Goal: describe the relevant legal framework	Goal: interpret the legal framework in a way which allows identifying a set of more specific objectives guiding the work of the PSM	Goal: specify the objectives using user needs	Goal: clarify which architectural measures (as part of a digital nudge) could be put in place (or could be made available to the users)
Swiss example. "The Swiss public service media shall contribute to ... free public opinion-forming through comprehensive, diverse and accurate information, especially regarding political, economic and social matters" (al. 4 let a Federal Act on Radio and Television)	Comprehensive, accurate information on political, economic and social realities	Educate me, update me	Links to information that is as factual as possible. - Opportunities to let users choose: thematic choices (choice list), choice of levels of depth (I'd like to understand in depth...) - Suggest links to specific events (e.g. preparation for an upcoming vote)
	Diversified information - different perspectives	Give me perspective, keep me on trend	Links to various readings / perspectives. Ethical risk: choice inherent in the idea of offering different perspectives (which perspectives?)
	Public education (educational programmes)	Educate me	Educational content (which content for which priorities?)

Table 3: Description of the legal framework, specific objectives, relevant user needs, and possible architectural measures for public service media.

goal are *culture* Inspire me, *monde* Divert me, *culture* Inspire me, *culture* Inspire me, *monde* Keep me on trend, *monde* Divert me, *economie* Keep me on trend, *monde* Inspire me;

2. "Citizen-perfectionist nudging" ($nudging = 1$, $favouriteTopic = None$): the target community is 15 and the suggested combinations are *culture* Inspire me, *monde* Divert me, *culture* Inspire me, *culture* Inspire me, *monde* Keep me on trend, *monde* Divert me, *economie* Keep me on trend, *monde* Inspire me, *sciences-tech* Divert me, *culture* Keep me on trend, *sciences-tech* Inspire me, *suisse* Divert me, *culture* Educate me, *culture* Inspire me. As expected, the first suggestions are the same because community 10 is closer to the given user than community 15;
3. "Random discovery" ($nudging = 0$): a random target community is identified (5) and the suggested article has topic *sciences-tech* and user need Divert me.

The second user has connection values randomly generated from 0 to 30 and is associated to community 14:

1. "Give-me-perspective nudging" ($nudging = 1$, $favouriteTopic = culture$): the target community is 0 and the suggested article for this scenario and user are *culture* Keep me on trend, *economie* Keep me on trend, *culture* Inspire me, meaning that the user is already close to the target community;
2. "Citizen-perfectionist nudging" ($nudging = 1$, $favouriteTopic = None$): the target community is 15 and the suggested articles have combinations *sport* Update me, *sciences-tech* Divert me, *culture* Keep me on trend, *culture* Inspire me, *sciences-tech* Inspire me, *sciences-tech* Keep me on trend, *culture* Inspire me, *culture Educate me*, *economie* Keep me on trend, *monde* Inspire me, *suisse* Divert me, *cultur*e Inspire me, *culture* Inspire me;
3. "Random discovery" ($nudging = 0$): a random target community is identified (6) and the suggested article has topic *monde* and user need Divert me.

As expected, different nudging scenarios lead to different articles being recommended, therefore proving the concept of the different recommendation approaches.

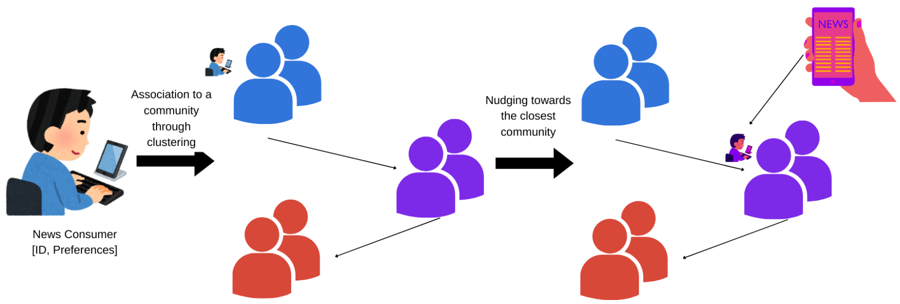

Fig. 5: Illustration of the proposed recommendation tool's workflow. Users are matched with communities based on their preferences. Following the implementation of nudging techniques, tailored articles are suggested to gradually engage users with their target community, contingent on the specific nudging applied.

5 Digital nudging and its normative challenges

The scenarios presented above can be approached as challenges of digital nudging [5]. The precedent section has shown the technical feasibility of this recommendation system. This section focuses on the normative challenges linked to these three nudging scenarios. It is organized into two main sections. The first section identifies and structures the relevant normative challenges. The second section outlines the three components of the legitimacy framework meant to address these challenges.
Here, legitimacy is understood as a normative concept rooted in the ethical alignment between a PSM's public service mandate, values, and commitments, rather than a sociological notion of perceived credibility. While legal norms (e.g., Swiss law) provide a broad operational margin, this space invites ethical scrutiny to ensure consistency between institutional goals and algorithmic practices.

5.1 Paternalism, perfectionism and manipulation

The use of algorithmic recommendation systems by PSMs can be understood as a form of digital nudging—non-coercive interventions in interface design that steer users toward certain choices without removing alternatives [6, 3]. In this context, PSMs influence user behavior not by restricting content, but by structuring its presentation (e.g., ordering or highlighting).
Two aspects of digital nudging are particularly relevant. First, architectural choices are unavoidable: there is no neutral way to present content. Even default settings like "most recent first" reflect a deliberate design decision. Second, while such choices are necessary, their ethical acceptability depends on how well they align with public service values. This is where the notion of legitimacy as consistency becomes essential. The following sections examine the ethical challenges of nudging through the lenses of paternalism, perfectionism, and manipulation.

Paternalism The challenge of paternalism concerns whether users are truly free and competent to choose what is best for themselves. An intervention is considered paternalistic when it aims to guide users toward outcomes deemed desirable, even without their consent or against their will [12].
Ethical debates on paternalism often center on its tension with freedom or autonomy. A key issue is identifying when such interference threatens individual liberty. In the context of Thaler and Sunstein's "libertarian paternalism," nudging is seen as non-coercive because it preserves choice [40, 2]. However, this depends on

a particular conception of freedom [10], and nudges still represent deliberate influence that may conflict with user preferences. Paternalistic approaches thus raise questions about procedural safeguards, such as transparency and informed consent—issues we revisit in relation to algorithmic recommendation systems.

This challenge is especially relevant for public service media (PSM). There is a clear power asymmetry between institutions and users, yet public bodies must uphold individual freedom and pluralism. Unlike private actors, PSM are bound by constitutional norms and must avoid promoting specific ethical or ideological convictions.

Perfectionism The use of nudges also raises the challenge of perfectionism, the idea that individuals should be guided toward a particular conception of the good (e.g., rationality, good citizenship) [15, 31]. Nudges, in this view, aim to steer behavior toward these ideals. This raises two main concerns: (1) how to justify the chosen standard of the good, and (2) whether it is legitimate for PSM to guide individuals toward it.

Perfectionist standards vary, some focus on achieving certain goods, others on developing human capacities like rationality. Even assuming a legitimate standard, it remains debatable whether PSM should aim to continually improve individuals according to that standard.

The main criticism of perfectionism comes from liberal neutrality: public institutions should not favor one ethical view over another. As Dworkin famously argued, the state shouldn't judge whether watching football is less worthy than visiting a museum [13]. Respect for ethical diversity requires restraint from public authorities.
For PSM, this raises a key tension: should they meet users as they are, respecting their interests and choices, or aim to improve them? Initiatives like Radio France's "touch of surprise" suggest a perfectionist intent, assuming that exposing users to unexpected content fosters personal development.

Manipulation At the intersection of paternalism and perfectionism, nudging raises concerns about gentle manipulation. Addressing this requires clarifying: (1) what counts as manipulation, and (2) why it is ethically problematic. Definitions of manipulation, as outlined by Noggle [27], fall into three main categories: Cognitive bypassing, i.e., influence that exploits cognitive biases and bypasses rational deliberation [38]; Deception-based, i.e., manipulation through hidden intentions or misleading methods, shaped by power dynamics [18]; Coercion-like, that is, not overtly deceptive, but manipulative through contextual or relational pressure, like emotional blackmail [14].
The ethical concern depends on the definition but typically involves two arguments. From one side, consequentialist: manipulation may cause harm, though defining "harm" is complex when outcomes seem beneficial (e.g., suicide prevention). From the other autonomy-based: manipulation undermines personal autonomy by obstructing critical reflection and disrespecting individuals' capacity for reasoned choice [11].
Notably, this critique does not require assuming humans are fully rational, it simply asserts that manipulation can deny them the chance to engage reflectively and autonomously.

5.2 A three-component legitimacy framework

The previous section has highlighted a set of normative challenges directly related to algorithmic recommendation systems as a digital nudge. This section outlines a legitimacy framework to address these challenges. It entails three components: the objectives (the "why"), the implementation modalities (the "how"), and an analysis of potential aggregated risks.

The objectives of the PSM A legitimate recommendation system must pursue legitimate objectives, especially within the framework of PSM, whose mandate is grounded in national and European legal norms. These norms reflect specific democratic values, such as universality, diversity, and independence, which vary by country. In Switzerland, for example, the legal framework emphasizes linguistic diversity, federalism, and direct democracy.
Our case study highlights how disagreements may arise in interpreting legal mandates, identifying user needs, or designing recommendation architectures. We propose a step-by-step approach to assess the legitimacy of recommendation objectives [32], illustrated through three nudging scenarios (Section 3): recommending Swiss political news aligned with Give me perspective; recommending content under Give me perspective without

a specific theme; suggesting content distant from the user's current interests.
These scenarios raise ethical concerns, particularly scenario 1, which combines paternalism and perfectionism. Scenario 2 also reflects a lighter form of perfectionism. PSM should acknowledge this by referencing the legislator's role in defining civic-oriented objectives. We distinguish between two forms of civic perfectionism:

- Light perfectionism, which promotes basic civic knowledge and skills (functional citizenship);
- Demanding perfectionism, which aims at an ideal, harder-to-define civic model.

Given the lack of baseline measures, light perfectionism is more legitimate and aligns with the public service mission. It supports promoting civic-relevant content (e.g., voting information) over purely entertainment content. Thus, discoverability in PSM should not be value-neutral. It should guide users toward content that enhances civic competence, aligning with public interest goals [8].

Modalities: Giving Control to Users Four key modalities address manipulation risks in algorithmic recommendation systems for PSM, based on respecting user autonomy:

- Transparent Information: Users must be clearly informed about the system's existence, operation, objectives, and implications in accessible language and format;
- Specific Operational Information: Users should understand how the system works, including data storage, use, deletion, architectural choices, and limitations, aligning with data protection laws;
- User Control: Users must control their choices, including data use, revising settings, opting out, and configuring features aligned with personal goals within PSM objectives;
- Informed Consent: Users should actively opt in rather than be automatically enrolled.

These four modalities bring recommendation systems closer to an ideal where users can autonomously configure their experience. This should inform system design, not just usage. For example, instead of tracking engagement, user preferences can be collected through multiple-choice questions, respecting privacy-by-design and data minimization principles. While the default system complies with data protection laws, users could opt into enhanced features involving limited tracking.
A key challenge for PSM is balancing user freedom with public service goals, e.g., should users be able to prioritize only entertainment and ignore civic content? Although configurability is ultimately shaped by design choices, systems can offer flexibility while reserving space for essential public service content (e.g., voting information). This approach reconciles user autonomy with the PSM mandate, rather than treating them as incompatible.

Aggregated Risks: going beyond the individual users The first two components of the legitimacy framework are user-focused, but must be complemented by an analysis of aggregated risks, those that extend beyond the individual level [24]. Inspired by the European AI Act, which mandates risk assessments for high-risk AI systems, this perspective emphasizes potential impacts on democracy, health, safety, and fundamental rights. While the Act may not directly apply to PSM, it illustrates the broader relevance of systemic risk analysis.
One example of a public-level aggregated risk concerns personalization and the loss of shared information spaces. As content becomes fragmented and asynchronously consumed—e.g., users watching only selected news clips, public discourse suffers due to the erosion of common informational ground [30], [21]. To counter this, PSM might prioritize promoting key democratic content to all users, overriding typical personalization.
A second example involves internal organizational risks. In editorial teams, performance metrics, though not commercially driven, can distort priorities if content fulfilling broad user needs is favored over more niche but important material [39]. This may harm team incentives and content diversity.
These examples are not exhaustive but illustrate why aggregated risk analysis must be integrated into the legitimacy framework, complementing individual-level concerns and guiding both content and architectural decisions [4].

6 Conclusion

The use of algorithmic recommendation systems by PSM raises important normative challenges. Based on real data, this study proposed three digital nudging scenarios built upon a novel recommendation approach using user needs. The first two focused on promoting civic content, aligning with a moderate perfectionist approach that supports civic engagement. The third, involving random content discovery, lacks sufficient justification and is less legitimate. Instead, a gradual, user-guided journey toward diverse content is recommended. The proposed legitimacy framework highlights the need for transparency to ensure user autonomy and maintain public trust. Transparent, explainable systems can set a standard for responsible digital media practices. Our recommendation tool, grounded in graph theory and machine learning, shows promise for enhancing content and tone diversity by incorporating user needs. We propose conducting real-world testing of our tool to evaluate changes in content consumption behavior. Until now, the tool has not been tested on the partner platform and lacks user testing.

In future research, we plan to incorporate more sophisticated machine-learning techniques, like Generative Adversarial Networks (GANs), to produce realistic synthetic data, increasing the accuracy of the results. Also, we aim to consider factors such as the timestamps and the sequence of articles visited by users. This will enable a more thorough understanding of individual user behaviors. Furthermore, it emerged from the digital report news 2024 [25] that: *'Most people don't want the news to be made more entertaining, but they do want more stories that provide more personal utility, help them connect with others, and give people a sense of hope'*. The focus on more varied user needs, such as Inspire me or Connect me, could therefore be considered in future applications, along with a careful evaluation of their ethical implications.

Acknowledgment

The authors would like to express their gratitude to the Initiative for Media Innovation (IMI) for providing the financial and institutional backing that made this research possible. We also extend our thanks to Radio Television Suisse (RTS) for their invaluable collaboration and support throughout this project.
The authors report there are no competing interests to declare.
Chat GPT was partially used for rephrasing.

References

1. Andersen, K., Shehata, A., Skovsgaard, M., Strömbäck, J.: Selective news avoidance: Consistency and temporality. Communication Research (01 2024). https://doi.org/10.1177/00936502231221689
2. Barton, A., Grüne-Yanoff, T.: From libertarian paternalism to nudging—and beyond. Review of Philosophy and Psychology **6**(3), 341–359 (2015). https://doi.org/10.1007/s13164-015-0268-x, https://doi.org/10.1007/s13164-015-0268-x
3. Bergram, K., Djokovic, M., Bezençon, V., Holzer, A.: The digital landscape of nudging: A systematic literature review of empirical research on digital nudges. CHI '22: Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems **62**(1-16) (2022)
4. Van den Bulck, H., Moe, H.: Public service media, universality and personalisation through algorithms: Mapping strategies and exploring dilemmas. Media, Culture & Society **40**(6), 875–892 (2017). https://doi.org/10.1177/0163443717734407, https://doi.org/10.1177/0163443717734407
5. Burri, M.: Nudging as a Tool of Media Policy: Understanding and Fostering Exposure Diversity in the Age of Digital Media, pp. 315–341. Springer (2015)
6. Caraban, A., Karapanos, E., Gonçalves, D., Campos, P.: 23 ways to nudge: A review of technology-mediated nudging in human-computer interaction. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems p. Paper 503 (2019). https://doi.org/10.1145/3290605.3300733
7. Carillon, K.: An algorithm for public service media: Constrained embedding of values into algorithmic design. Sage Journals (2024), https://journals.sagepub.com/doi/10.1177/27523543241290976?icid=int.sj-abstract.citing-articles.3, accessed May 8, 2025
8. Committee, o.M.: Recommendation of the committee of ministers to member states on principles for media and communication governance (adopted by the committee of ministers on 6 april 2022 at the 1431st meeting of the ministers' deputies). https://search.coe.int/cm?i=0900001680a61712 (2022)
9. Diboune, A., Slimani, H., Nacer, H., Beghdad Bey, K.: A comprehensive survey on community detection methods and applications in complex information networks. Social Network Analysis and Mining **14**(1), 93 (2024). https://doi.org/10.1007/s13278-024-01246-5
10. Dowding, K., Oprea, A.: Nudges, regulations and liberty. British Journal of Political Science **53**(1), 204–220 (2023). https://doi.org/10.1017/S0007123421000685, https://www.cambridge.org/core/product/04D503466F0B0FC5870266D57918DE4F

11. Dworkin, G.: The nature of autonomy. Nordic Journal of Studies in Educational Policy **2015**(2), 1–14 (2015-01-01). https://doi.org/10.3402/nstep.v1.28479
12. Dworkin, G.: Paternalism. Stanford Encyclopedia of Philosophy (2020)
13. Dworkin, R.: A Matter of Principle. Harvard University Press (1995)
14. Faden, R.R., Beauchamp, T.L., King, N.M.P.: A History and Theory of Informed Consent. Oxford University Press (1986)
15. Foot, P.: Natural Goodness. Oxford University Press (2003)
16. Fortunato, S.: Community detection in graphs. Physics Reports **486**(3), 75–174 (2010). https://doi.org/https://doi.org/10.1016/j.physrep.2009.11.002
17. Frisch, L., Veil, S.: Radio france instaure un nouvel algorithme associant performance et découverte (2022), https://www.freelance-informatique.fr/actualites/radiofrance-nouvelalgorithme-performancedecouverte, accessed May 8, 2025
18. Goodin, R.E.: Manipulatory Politics. Yale University Press (1980)
19. Heitz, L., Lischka, J., Birrer, A., Paudel, B., Tolmeijer, S., Laugwitz, L., Bernstein, A.: Benefits of diverse news recommendations for democracy: A user study. Digital Journalism **10**(10), 1710–1730 (2022). https://doi.org/10.1080/21670811.2021.2021804
20. Jones, E.: Inform, educate, entertain… and recommend. Ada Lovelace Institute (2022)
21. Just, N., Latzer, M.: Governance by algorithms: Reality construction by algorithmic selection on the internet. Media, Culture & Society **39**(2), 238–258 (2017-03-01). https://doi.org/10.1177/0163443716643157
22. Lancichinetti, A., Fortunato, S.: Community detection algorithms: A comparative analysis. Phys. Rev. E **80**, 056117 (2009). https://doi.org/10.1103/PhysRevE.80.056117
23. Milano, S., Institute, A.L.: How public service media organizations can create a responsible approach to algorithmic recommendations. Tech Xplore (2022), https://techxplore.com/news/2022-11-media-responsible-approach-algorithmic.html, accessed May 8, 2025
24. Milano, S., Taddeo, M., Floridi, L.: Recommender systems and their ethical challenges. AI & SOCIETY **35**(4), 957–967 (2020). https://doi.org/10.1007/s00146-020-00950-y, https://doi.org/10.1007/s00146-020-00950-y
25. Newman, N., Fletcher, R., Robertson, C.T., Arguedas, A.R., Nielsen, R.K.: Reuters institute digital news report 2024. Reuters Institute (2024)
26. Ni, T., Zhu, R., Krever, R.: Responses to news overload in a non-partisan environment: News avoidance in china. SAGE Open **13** (2023). https://doi.org/10.1177/21582440231184864
27. Noggle, R.: The ethics of manipulation. Stanford Encyclopedia of Philosophy (2022)
28. Olsson, T., Huhtamäki, J., Kärkkäinen, H.: Directions for professional social matching systems. Communications of the ACM **63**(2), 60–69 (2020)
29. Pattanayak, H.S., Verma, H.K., Sangal, A.L.: Community detection metrics and algorithms in social networks. In: 2018 First International Conference on Secure Cyber Computing and Communication (ICSCCC). pp. 483–489 (2018). https://doi.org/10.1109/ICSCCC.2018.8703215
30. Radio, W., Council, T.: Public broadcasting: why? how? UNESCO (2001)
31. Rawls, J.: A Theory of Justice. Belknap Press of Harvard University Press (1971)
32. Rochel, J.: Recommandation algorithmique et nudging: un modèle de légitimité pour le service public audio-visuel. Medialex pp. 1–20 (2024)
33. Rolandsson, T., Widholm, A., Rahm-Skågeby, J.: Managing public service: The harmonization of datafication and managerialism in the development of a news-sorting algorithm. Digital Journalism **10**(10), 1691–1709 (2022-11-26). https://doi.org/10.1080/21670811.2022.2119151
34. Schmidt, A.T., Engelen, B.: The ethics of nudging: An overview. Philosophy Compass **15**(4), 1–13 (2020). https://doi.org/https://doi.org/10.1111/phc3.12658, https://onlinelibrary.wiley.com/doi/abs/10.1111/phc3.12658
35. Shishkin, D.: User needs in content publishing: the slide that started it all, five years later — linkedin.com. https://www.linkedin.com/pulse/user-needs-content-publishing-slide-started-all-five-years-shishkin/ (2022)
36. Smartocto: User needs model for news. https://smartocto.com/research/userneeds/ (2023)
37. Sørensen, J., Hutchinson, J.: Algorithms and Public Service Media, pp. 91–106. Re-Visionary-Interpretations of the Public Enterprise (RIPE), Nordicom (Apr 2018), http://ripeat2016.org, invited chapter; RIPE@2016 ; Conference date: 21-09-2016 Through 24-09-2016
38. Susser, D., Roessler, B., Nissenbaum, H.: Technology, autonomy, and manipulation. Internet Policy Review **8**(2) (2019)
39. Sørensen, J.K.: Public service media, diversity and algorithmic recommendation: Tensions between editorial principles and algorithms in european PSM organizations. CEUR Workshop Proceedings pp. 6–11 (2019)
40. Thaler, R., Sunstein, C.R.: Nudge : Improving decisions about health, wealth, and happiness. Yale University Press, New Haven (2008)
41. Traag, V.A., Waltman, L., van Eck, N.J.: From louvain to leiden: guaranteeing well-connected communities. Scientific Reports **9**(1) (2019). https://doi.org/10.1038/s41598-019-41695-z

42. Wang, Y., Hao, S., Wang, X., Zhuang, X.: Enhancing graph topology and clustering quality: A modularity-guided approach (2024)