

Evaluating Image Trust Labels in a News Recommender System

Svenja Lys Forstner^{1,†}, Yelyzaveta Lysova^{1,*,†}, Alain D. Starke¹ and Christoph Trattner¹

¹MediaFutures, University of Bergen, Lars Hilles Gate 30, 5008 Bergen, Norway

Abstract

Rising user concerns about online misinformation and the spread of AI-generated visual content underscore the need for better ways to verify image authenticity. Image provenance labels are a proposed solution, aiming to help users assess the veracity of digital images. The Coalition for Content Provenance and Authenticity (C2PA), for instance, can disclose image provenance (i.e., origin or source details) to users in the form of labels that describe the image’s metadata. However, little is known about whether users engage with or understand such labels, especially in news recommender contexts. In this paper, we introduce an alternative ‘Image Trust Score’ label, inspired by the front-of-package Nutri-Score label, and experimentally evaluate its effectiveness in a personalized news setting. We present the results of a four-condition (no-label baseline, C2PA label, black-and-white and colored Image Trust Score) between-subjects study ($N = 202$) in which participants selected news articles (with or without labels), reporting on label comprehension and trust. While image trust and article selection were not significantly affected, all labels increased article trust. The Image Trust Score was perceived as more understandable and appealing than the C2PA label, though many participants misinterpreted the labels’ meaning. Our findings highlight the need for clearer and more intuitive provenance label design.

Keywords

Recommender Systems, News Recommender Systems, Trust, Image Provenance, C2PA, Provenance Labels, User Trust, Nudging

1. Introduction

Recent advances in generative AI have made the creation and manipulation of visual content more accessible [1, 2], which is fueling the spread of misinformation in digital media [3, 4, 5]. This development has changed how trust in online media content is formed [6], as people may be more skeptical or simply lack the skills to detect or understand the authenticity of online content [7, 8].

News consumers can be supported in avoiding misinformation. One approach is to visually highlight the extent to which information and media content are verified [9]. A promising approach involves using visual indicators such as image provenance labels [10], which provide details about an image’s origin and editing history. The Coalition for Content Provenance and Authenticity (C2PA) enables this by embedding digitally signed metadata in the form of ‘Content Credentials’ across the entire news production chain [11].

Effective communication of C2PA metadata to end users is a challenge. While outlets like the BBC have trialed the framework in select publications [12], adoption remains limited and user awareness is low. Despite growing interest in provenance indicators, research on their impact on user trust and engagement—particularly within news recommender systems—is scarce [13, 12, 10].

To address this gap, we examine how disclosing C2PA metadata via image provenance labels affects users in a news recommender system. We compare two label types (see Figure 1): the established C2PA label, which indicates metadata availability, and a newly designed “Image Trust Score” label,

Proceedings of the 13th International Workshop on News Recommendation and Analytics (INRA 2025), co-located with the 19th ACM Conference on Recommender Systems (RecSys 2025), September 22–26, 2025, Prague, Czech Republic.

*Corresponding author.

[†]These authors contributed equally.

✉ svenja.forstner@uib.no (S. L. Forstner); y.lysova@student.uib.no (Y. Lysova); alain.starke@uib.no (A. D. Starke); christoph.trattner@uib.no (C. Trattner)

ORCID 0009-0003-9738-0803 (S. L. Forstner); 0009-0006-9605-8291 (Y. Lysova); 0000-0002-9873-8016 (A. D. Starke); 0000-0002-1193-0508 (C. Trattner)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

which is inspired by nutrition panels like the Nutri-Score [14] and uses a letter-grade rating to convey trustworthiness more intuitively. Given the novelty of C2PA technology in news recommendation, we assess how labeled articles influence user engagement, and further compare the labels in terms of perceived trust, comprehensibility, and intuitive design [15], guided by the following research questions:

- **RQ1:** To what extent do different types of image provenance labels influence readers’ trust perceptions of news articles and images in a news recommender system?
- **RQ2a:** How accurately do readers interpret different types of image provenance labels, and how is this influenced by their prior familiarity with similar labels?
- **RQ2b:** How do readers’ self-assessments of their understanding of image provenance labels compare to their actual comprehension?

2. Background

2.1. Nudging and Recommender Systems

The effects of using labels in recommender interfaces can be explained through Nudge theory. A nudge is a design aspect of an interface (i.e., the ‘choice architecture’), which leads to predictable changes in user behavior [16, 17], such as through what option is chosen. Predictable in this definition refers to the nudge tapping into a cognitive aspect on the user’s end that elicits a specific response [17]. For example, if an option in a recommender interface is pre-selected [18], this is referred to as a default option and is more likely to be chosen because of the status quo bias [19, 20].

Labels can be considered akin to the use of explanations in recommender systems [21, 22], increasing the salience of a specific aspect of what is recommended to an end user [23]. For example, explanations or nutrition labels in a recipe recommender system can emphasize the healthiness of different options [24, 25, 26], nudging users to value health attributes higher in their recipe choices [27, 18]. By analogy to nutrition labels, we expect that image provenance labels emphasize the importance of verification practices, acting as trust-based nudges.

2.2. Provenance and Nutrition Labels

Recent initiatives highlight how labels communicate credibility. C2PA is a collaborative, industry-driven standard for embedding cryptographically signed “Content Credentials” into images and other media, recording the source, editing steps, and responsible parties of visual content [11]. Pilots like IPTC’s Origin Verified Publisher certificates (e.g., BBC, CBC), Deutsche Welle’s on-page labels, and France Télévisions’ newsroom integration have demonstrated a proof-of-concept [28, 29, 30].

Similarly, front-of-pack nutrition labels (e.g., Nutri-Score, Traffic Light, Guideline Daily Amounts) offer consumers a quick, at-a-glance assessment of a food product’s healthfulness using color codes, symbols, and numeric or letter grades to communicate complex nutritional data [14, 31, 32]. Studies show that well-designed nutrition labels can improve consumer understanding and influence purchasing decisions, making them a compelling model for conveying provenance data in other domains like news media [14].

2.3. Contributions

This paper makes three key contributions to the design and evaluation of provenance labels in news recommender systems. First, we introduce the *Image Trust Score* label, inspired by front-of-pack nutrition panels, using icons, color, scores, and text to communicate image provenance. Second, we integrate both the C2PA and Image Trust Score labels into a News Recommender prototype, enabling direct comparison of user engagement with labeled and unlabeled images. Third, we report findings from a user study evaluating the comprehensibility and intuitiveness of both labels, addressing the impact on user trust and selection behavior.

3. Methods

3.1. Dataset

We created a dataset of news articles from major international English news outlets [5]: nytimes.com, reuters.com, foxnews.com, cbsnews.com, washingtonpost.com, and cnn.com. To avoid potential bias from participants' opinions of or trust in these outlets, the original sources were not displayed during the experiment. Instead, we assigned articles fictional newspaper names, which appeared as the source in the detailed C2PA label (see Figure 1a). For each of the categories Politics, Sports, Technology, and Entertainment, 15 articles were sampled. Article selection was performed via the NewsCatcher API [33], and only items published within the five days preceding the study were included to ensure topical currency. The study prototype, R scripts and analysis plots are available [here](#).

3.2. Participants

Participants were recruited through the research platform Prolific [34]. Of the final sample that completed the study and passed all attention checks ($N = 202$), 104 participants were residents of the United Kingdom and 98 were from the United States. The average study completion time was 10:45 minutes, with a remuneration of £8.67–£11.64 per hour. In terms of gender, 51.0% identified as female, 47.5% as male, 1.0% as non-binary, and 0.5% preferred not to say, providing a representative sample. Participants ranged in age from 16 to over 65, with the largest cohort aged 25–34 years (29.7%) and the smallest aged 65 or older (3.0%).

3.3. Research Design

We employed a between-subjects design, with participants randomly assigned to one of four conditions. The baseline group saw news articles without labels. The other three groups viewed articles where the cover image was accompanied by either a C2PA label (Figures 1a, 1d) or an *Image Trust Score* label—specifically designed for this study—in color (C-ITS, Figures 1b, 1e) or black and white (BW-ITS, Figures 1c, 1f). Since the C2PA label presents metadata without interpretation, the Image Trust Score was designed to summarize it using a simplified format and an A–E letter grade, providing a clearer sense of metadata quality. To minimize confounding from grade interpretation, only the top score (“A”) was shown. Each participant completed three rounds of news article selection, viewing the full article after each choice and then selecting the next from a set of four recommended articles, two of which displayed a label. The baseline group received one attention check, while labeled conditions included two due to additional questionnaire items.

3.4. Procedure


Following informed consent and demographic questions, participants completed a questionnaire about their news consumption habits and general trust in news media. They then entered the main experimental task, where they were shown a grid of four news previews, each corresponding to one of the selected news topics (see Section 3.1). Depending on the assigned condition, each preview displayed a *simple* label (see Figures 1d–1f), with the exception of the baseline group which saw no labels.

After their initial selection, participants completed three rounds of article interaction. In each round, the selected article was presented in full, accompanied by its cover image, a condition-specific *detailed*



Figure 1: Label set used in the study. (a)–(c): *Detailed* C2PA, C-ITS, and BW-ITS labels presented with full-size articles, respectively; (d)–(f): Corresponding *simple* labels presented with recommended articles.

Taiwan adds China's Huawei, SMIC to export control list



What is the Image Trust Score?

This score rates how transparent and trustworthy the image is — based on 7 criteria:

- A: High transparency and verified information
- B-C: Some elements missing or unclear
- D-E: Poor or no transparency indicators

Sofie Zhang | July 12, 2025

Taiwan has added China's Huawei Technologies (HWT.UL) and Semiconductor Manufacturing International Corp (SMIC) (0981.HK), opens new tab to its export control list, which includes other proscribed organisations like the Taliban and al Qaeda. Inclusion on the economy ministry's trade administration's strategic high-tech commodities

(a)

Recommended


Ocado, Bon Preu to build new fulfilment centre in Catalonia



Apple eyes using AI to design its chips, technology executive says



Anduril, Rheinmetall partner to build military drones for Europe



Swedish military joins Telia, Ericsson to boost defense tech



(b)

Figure 2: Study user interface shown in each round of article interaction and selection (news images are blurred in this example): (a) Example of full article including the cover image with a *detailed* label and the expanded label explanation; (b) A recommended articles showing the *simple* label version.

label (see Figures 1a-1c), and a label explanation accessible via the label itself or an adjacent question mark icon (see Figure 2a). Participants were then asked to select from four recommended articles, two of which displayed the *simple* version of the respective condition's label below the cover image (see Figure 2b). The placement of labels among the recommended articles was randomized.

After each article selection, prior to viewing the full article, participants completed a brief mid-round questionnaire on their decision-making and trust in the selected article. Upon completing all three rounds, participants answered a final questionnaire focused on their perceptions of the labels, including an additional assessment of their grading scheme understanding in the C-ITS and BW-ITS conditions. The session concluded with a debriefing and summary of the study's aims.

3.5. Measures

Pre-task measures included self-reports of participants' news consumption frequency, primary news platform type, devices used for news consumption, and most-accessed news sources (with source options based on U.S. and U.K. popularity [5]). Participants further assessed their general trust in news media with a 7-point Likert scale item adapted from Strömbäck et al. [7], "I generally trust information from the news media in my country".

For each round of the main task, we recorded participants' article choices and whether they accessed the label explanation. After each article selection, participants indicated which news article preview elements it was influenced by (multiple-choice: *image label*, *image*, *headline*, *topic*, *other*), and rated the selected article regarding trust ("I trust this article most among all 4 recommended articles.") and perceived image trustworthiness ("I think the cover image in this article is trustworthy") on a 7-point Likert scale.

The post-task questionnaire assessed participants' interpretations of the labels. This multiple-choice item included both correct (*How much verified information about the creation and edits of the picture is available*, *How much of the picture is real/fake*, *How trustworthy the image is*) and incorrect (*How trustworthy the article is*, *How much readers liked the article*) options. Respondents were further asked to indicate what information they expected to see when clicking on the label (multiple-choice), and to rate the labels on several 7-point Likert items: *immediate understanding*, *visual appeal*, *usefulness for selection*, *informativeness*, *reassurance about image trustworthiness*, *support for evaluating image trust*, and *preference for more articles to display such labels*. Participants in C-ITS or BW-ITS conditions responded to an additional question on what they believed the label's displayed grade "A" was based on

(multiple-choice).

All participants reported their familiarity with both trust labels on social media and with the Nutri-Score food label (Yes/No/Not Sure). They also rated their confidence in identifying trustworthy news after completing the study (drop-down selection), and could provide further comments in an open-text field.

4. Results

4.1. Trust in Article Images and Articles

To compare perceived trust in images and articles across experimental conditions, we analyzed trust ratings from the Likert-scale items in the mid-round questionnaire, using one-way ANOVAs and applied Tukey HSD post-hoc tests for pairwise comparisons. C-ITS received the highest mean image trust rating ($M = 5.20$), followed by BW-ITS ($M = 5.14$), C2PA ($M = 4.93$), and No Label with the lowest rating ($M = 4.81$). However, Tukey HSD post-hoc tests indicated that none of the pairwise differences reached statistical significance ($p > .05$). In contrast, article trust ratings (see Figure 3) revealed significant effects: all three labeling conditions produced higher article trust than the No Label condition. Specifically, trust was significantly higher with C-ITS ($p < 0.001$), BW-ITS ($p = 0.012$), and C2PA ($p = 0.015$) when compared to the baseline condition, with no significant differences among the labeling conditions ($p > 0.1$). Thus, the presence of labels increased perceived article trust, though the specific type of label did not make a significant difference.




Figure 3: (a) Mean article trust in recommended articles by label condition; (b) Tukey HSD pairwise comparisons

4.2. Comparative Evaluation of Label Perception

Participants' perceptions of the image labels were analyzed using Likert-scale items from the final questionnaire. Both C-ITS and BW-ITS conditions consistently received higher ratings than C2PA across all dimensions. C-ITS scored highest for immediate understanding ($M = 4.60$), visual appeal ($M = 5.04$), and informativeness ($M = 4.94$). BW-ITS was rated highest regarding support for evaluating image trust ($M = 4.86$) and reassurance about image trustworthiness ($M = 4.82$). Both Image Trust Score conditions were rated significantly higher than C2PA in usefulness for selection (C-ITS $M = 4.38$, BW-ITS $M = 4.40$, C2PA $M = 3.82$), and in preference for more articles to display such labels (C-ITS $M = 5.29$, BW-ITS $M = 5.34$, C2PA $M = 4.98$). Notably, preference for wider label adoption was the highest-rated item across all conditions.

4.3. Label Interpretation and Understanding

Data for this analysis were drawn from multiple-choice responses regarding assumed representation of the label and the Likert-scale items across all labeled conditions in the final questionnaire. The

analysis showed that a number of participants misinterpreted the label’s meaning: 44% selected “*How trustworthy the article is*”, despite the label referring only to the image. A small number of participants (8%) selected the second incorrect answer option “*How much readers like the article*”.

Self-reported label understanding was highest and most consistent in the C-ITS condition (Mdn ≈ 5), followed by the BW-ITS (Mdn ≈ 4.5), and lowest for C2PA (Mdn ≈ 4). Clicks on the label or question mark icon for an explanation were low across all conditions (C-ITS 2.0%, BW-ITS 6.0%, C2PA 4.0%), suggesting limited spontaneous interest in accessing additional label information.

4.4. Familiarity and understanding

To assess the impact of prior exposure on label comprehension, we compared self-reported familiarity with social-media trust labels and the Nutri-Score to label comprehension scores by condition (see Figure 4). Across conditions, participants familiar with these indicators reported the highest understanding (e.g., C-ITS: $M \approx 4.82$ for trust label familiarity and $M \approx 4.74$ for Nutri-Score familiarity).




Figure 4: Self-reported understanding of labels (1 = strongly disagree, 7 = strongly agree) by prior trust labels familiarity (Yes/No/Not Sure) across label conditions

5. Discussion

Our results offer key insights for designing image provenance labels in News Recommender Systems. First, effective labels must be accessible, visually clear, and supported by concise explanations. Since labels like C2PA are still unfamiliar to most users, simply placing a label is insufficient without guidance. While labels did not significantly affect article choice or image trust, all three increased trust in the recommended article, suggesting a *halo effect* where image credibility influences perceptions of the article itself.

Nearly half of participants misinterpreted the labels as indicators of article trust, reinforcing this effect. The simplified C2PA label was hardest to interpret, especially for users unfamiliar with trust cues or rating systems like Nutri-Score, while color-coded designs improved comprehension. Few participants accessed further explanations of the labels. This low engagement may reflect limited salience or effort required for an extra click, or the perception that the label’s meaning was already clear. These findings highlight the importance of integrating brief, prominent explanations directly within the label interface, minimizing user effort.

A comparison of self-reported and actual understanding revealed a gap: participants exposed to C-ITS performed better in understanding and felt confident, whereas C2PA users often overestimated their comprehension. This misalignment underscores the need for design elements that help match readers’ label perceptions with their actual understanding. Repeated exposure to provenance labels,

particularly those clearly explained and visually distinct, may further enhance both comprehension and trust over time. Ultimately, label effectiveness may depend on perceived relevance and clarity to the user.

5.1. Limitations

Our study has some limitations. Only 6.9% of participants reported a news aggregator as a platform used for news consumption, while 44.6% cites news websites as their primary source. As the study simulated an interface similar to a news aggregator, these news habits may have influenced participants' interactions with the study. Second, our sample included only participants from the United Kingdom and the United States. Results may differ in other regions [5]. Although most participants stated to primarily access news via mobile devices, the study was conducted exclusively on computer screens, which may have affected interaction with the study. Finally, the C2PA condition in this study relied on a simplified, static version of the label. This could have influenced label perception and trust ratings for both images and articles [10].

5.2. Future Work

Future research could test similar setups with improved explanatory cues or visual affordances. In addition, it would be valuable to test these labels in different contexts, such as on social media platforms, where baseline trust and exposure to disinformation vary. This would help determine whether label effectiveness is context-sensitive and whether certain formats are better suited to more informal or fast-paced media environments.

6. Conclusion

This paper introduced the *Image Trust Score*, a provenance label inspired by front-of-pack nutrition designs, and compared it with a simplified C2PA label in a news recommender prototype. The study showed that while image trust itself did not change significantly, all labels increased trust in the associated article, suggesting a *halo effect*. The Image Trust Score was more intuitive and better understood than the C2PA label, which many users misinterpreted. We also observed a confidence-competence gap: participants often felt that they understood the labels, even when their interpretation was partly incorrect. These findings highlight the need for visually clear and self-explanatory label designs, supported by embedded explanations.

Acknowledgments

This work was supported by the Research Council of Norway with funding to MediaFutures: Research Centre for Responsible Media Technology and Innovation, through the Centre for Research-based Innovation scheme, project number 309339.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT-4 in order to: Paraphrase and reword. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] J. Hwang, S. Oh, A brief survey of watermarks in generative ai, in: 2023 14th International Conference on Information and Communication Technology Convergence (ICTC), 2023, pp. 1157–1160. doi:10.1109/ICTC58733.2023.10392465.
- [2] J. Langguth, K. Pogorelov, S. Brenner, P. Filkuková, D. T. Schroeder, Don't trust your eyes: image manipulation in the age of deepfakes, *Frontiers in Communication* 6 (2021). doi:<https://doi.org/10.3389/fcomm.2021.632317>.
- [3] E. Cavaciuti-Wishart, S. Heading, K. Kohler, S. Zahidi, Global Risks Report 2024, Insight Report, World Economic Forum, 2024. URL: <https://www.weforum.org/publications/global-risks-report-2024/>.
- [4] Survey on the Impact of Online Disinformation and Hate Speech, Study, UNESCO and Ipsos, 2023. URL: https://www.unesco.org/sites/default/files/medias/fichiers/2023/11/unesco_ipsos_survey.pdf.
- [5] N. Newman, A. R. Arguedas, C. T. Robertson, R. K. Nielsen, R. Fletcher, Digital News Report 2025, Research Report, Reuters Institute for the Study of Journalism, 2025. doi:10.60625/risj-8qqf-jt36.
- [6] S. Sahebi, P. Formosa, The ai-mediated communication dilemma: epistemic trust, social media, and the challenge of generative artificial intelligence, *Synthese* 205 (2025). doi:10.1007/s11229-025-04963-2.
- [7] J. Strömbäck, Y. Tsfati, H. Boomgaarden, A. Damstra, E. Lindgren, R. Vliegenthart, T. Lindholm, News media trust and its impact on media use: Toward a framework for future research, *Annals of the International Communication Association* 44 (2020) 139–156. URL: <https://doi.org/10.1080/23808985.2020.1755338>. doi:10.1080/23808985.2020.1755338.
- [8] J. Swart, Tactics of news literacy: How young people access, evaluate, and engage with news on social media, *New Media & Society* 25 (2023) 505–521. URL: <https://doi.org/10.1177/14614448211011447>. doi:10.1177/14614448211011447, first published online May 2, 2021.
- [9] T. J. Thomson, D. Angus, P. Dootson, E. Hurcombe, A. Smith, Visual mis/disinformation in journalism and public communications: Current verification practices, challenges, and future opportunities, *Journalism Practice* 16 (2020) 938–962. URL: <https://doi.org/10.1080/17512786.2020.1832139>. doi:10.1080/17512786.2020.1832139.
- [10] C. Trattner, S. L. Forstner, A. D. Starke, E. Knudsen, C2pa provenance labels increase trust in news platforms across western countries, 2025. doi:https://doi.org/10.31219/osf.io/pdhaz_v1, preprint OSF.
- [11] C2PA Steering Group, About C2PA, <https://c2pa.org/about/>, 2025. Accessed July 8, 2025.
- [12] L. Monday, L. Strappelli, Does provenance build trust?, 2024. URL: <https://www.bbc.co.uk/rdnewslabs/news/does-provenance-build-trust>, accessed July 17, 2025.
- [13] K. J. K. Feng, N. Ritchie, P. Blumenthal, A. Parsons, A. X. Zhang, Examining the impact of provenance-enabled media on trust and accuracy perceptions, *Proceedings of the ACM on Human-Computer Interaction* 7 (2023) 270:1–270:42. URL: <https://doi.org/10.1145/3610061>. doi:10.1145/3610061.
- [14] L. Dréano-Trécant, M. Egnell, S. Hercberg, P. Galan, J. Soudon, M. Fialon, M. Touvier, E. Kesse-Guyot, C. Julia, Performance of the front-of-pack nutrition label nutri-score to discriminate the nutritional quality of food products: A comparative study across 8 european countries, *Nutrients* 12 (2020) 1303. URL: <https://doi.org/10.3390/nu12051303>. doi:10.3390/nu12051303, accessed July 7, 2025.
- [15] S. Raza, C. Ding, News recommender system: A review of recent progress, challenges, and opportunities, *Artificial Intelligence Review* 55 (2022) 749–800. URL: <https://link.springer.com/article/10.1007/s10462-021-10043-x>. doi:10.1007/s10462-021-10043-x, accessed July 7, 2025.
- [16] E. J. Johnson, S. B. Shu, B. G. C. Dellaert, C. Fox, D. G. Goldstein, G. Häubl, R. P. Larrick, J. W. Payne, E. Peters, D. Schkade, B. Wansink, E. U. Weber, Beyond nudges: Tools of a choice architecture, *Marketing Letters* 23 (2012) 487–504. URL: <http://www.jstor.org/stable/23259227>. doi:10.1007/s11002-012-9186-1.

- [17] R. H. Thaler, C. R. Sunstein, *Nudge: The final edition*, Penguin, 2021.
- [18] A. D. Starke, M. C. Willemsen, Psychologically informed design of energy recommender systems: Are nudges still effective in tailored choice environments?, in: B. Ferwerda, M. Graus, P. Germanakos, M. Tkalčič (Eds.), *A Human-Centered Perspective of Intelligent Personalized Environments and Systems*, Human-Computer Interaction Series, Springer, Cham, 2024, pp. 221–259. URL: https://doi.org/10.1007/978-3-031-55109-3_9. doi:10.1007/978-3-031-55109-3_9.
- [19] D. G. Goldstein, E. J. Johnson, A. Herrmann, M. Heitmann, Nudge your customers toward better choices, *Harvard Business Review* 86 (2008) 99–105. URL: <https://hbr.org/2008/12/nudge-your-customers-toward-better-choices>.
- [20] E. J. Johnson, D. G. Goldstein, Do defaults save lives?, 2003. URL: <https://ssrn.com/abstract=1324774>. doi:10.1126/science.1091721.
- [21] I. Nunes, D. Jannach, A systematic review and taxonomy of explanations in decision support and recommender systems, *User Modeling and User-Adapted Interaction* 27 (2017) 393–444. URL: <https://doi.org/10.1007/s11257-017-9195-0>. doi:10.1007/s11257-017-9195-0.
- [22] N. Tintarev, J. Masthoff, Designing and evaluating explanations for recommender systems, in: F. Ricci, L. Rokach, B. Shapira, P. B. Kantor (Eds.), *Recommender Systems Handbook*, Springer US, 2011, pp. 479–510. URL: https://doi.org/10.1007/978-0-387-85820-3_15. doi:10.1007/978-0-387-85820-3_15.
- [23] M. Jesse, D. Jannach, Digital nudging with recommender systems: Survey and future directions, *Computers in Human Behavior Reports* 3 (2021) 100052. URL: <https://doi.org/10.1016/j.chbr.2020.100052>. doi:10.1016/j.chbr.2020.100052.
- [24] A. E. Majjodi, A. D. Starke, C. Trattner, Nudging towards health? examining the merits of nutrition labels and personalization in a recipe recommender system, in: *Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization, UMAP '22*, Association for Computing Machinery, New York, NY, USA, 2022, pp. 48–56. URL: <https://doi.org/10.1145/3503252.3531312>. doi:10.1145/3503252.3531312.
- [25] A. Starke, E. Asotic, C. Trattner, “serving each user”: Supporting different eating goals through a multi-list recommender interface, in: *Proceedings of the 15th ACM Conference on Recommender Systems (RecSys '21)*, Association for Computing Machinery, New York, NY, USA, 2021, pp. 124–132. doi:10.1145/3460231.3474232.
- [26] A. D. Starke, A. El Majjodi, C. Trattner, Boosting health? examining the role of nutrition labels and preference elicitation methods in food recommendation, 2022. URL: <https://hdl.handle.net/11250/3040188>.
- [27] A. D. Starke, E. Asotic, C. Trattner, E. J. Van Loo, Examining the user evaluation of multi-list recommender interfaces in the context of healthy recipe choices, *ACM Transactions on Recommender Systems* 1 (2023) 1–31. doi:10.1145/3581930.
- [28] International Press Telecommunications Council, IPTC to create a C2PA-compatible list of Verified News Publishers, including BBC and CBC, <https://iptc.org/news/iptc-c2pa-verified-news-publishers/>, 2024. Accessed July 5, 2025.
- [29] Deutsche Welle Innovation, Pilot Testing of C2PA Labels in News Workflows, Deutsche Welle Innovation (2025). URL: <https://innovation.dw.com/articles/digital-seals-trust-news-media>, accessed July 7, 2025.
- [30] France Télévisions & Dalet, Proof-of-Concept: C2PA Integration in Newsroom Systems, <https://www.dalet.com/blog/trust-news-france-televvisions-combat-misinformation/>, 2025. Accessed July 7, 2025.
- [31] Food Standards Agency (FSA), Check the Label: Front-of-Pack Traffic Light Nutrition Labelling, <https://www.food.gov.uk/safety-hygiene/check-the-label>, 2021. Accessed July 7, 2025.
- [32] M. Rayner, P. Scarborough, C. Williams, The origin of guideline daily amounts and the food-based “rules of thumb” for front-of-pack nutrition labelling in the uk, *Public Health Nutrition* 7 (2004) 549–556. URL: <https://pubmed.ncbi.nlm.nih.gov/15153261/>. doi:10.1079/PHN2003552, accessed July 7, 2025.
- [33] NewsCatcher, Newscatcher api for news, ??? URL: <https://www.newscatcherapi.com/>

product-overview, accessed July 4, 2025.

[34] Prolific, Easily collect high-quality data from real people, <https://www.prolific.com>, 2025.