Chapter 4
Boosting Consumers: Algorithm-Supported Decision-Making under Uncertainty to (Learn to) Navigate Algorithm-Based Decision Environments

Felix G. Rebitschek

Human choice, for example, in decisions to consume goods or services or to participate in organizations and events, depends on seeking quality assured, objectively required, and subjectively needed information (Fritz & Thiess, 1986). Whereas in pre-digital days searching for information required substantial efforts, digitalization has improved information accessibility and facilitated consumers' information searches. Individual consumers, however, nowadays face comprehensive sets of information and more offers about products and services than they have the resources to navigate (Lee & Lee, 2004). Selecting information and preventing information overload have become major challenges for preparing consumer decisions (Glückler & Sánchez-Hernández, 2014).

Given the complexity and dynamics, information selection is a decision-problem under uncertainty. Distinct from problems of risk, problems of uncertainty are characterized by a lack of reliable evidence on choice options, the potential consequences of pursuing or not pursuing those options, and the probabilities of those consequences setting in (Knight, 1921). In contrast to non-reducible aleatory uncertainty (e.g., the next coin flip), these are problems of epistemic uncertainty that actors need to use knowledge to reduce.

Opposite the decision-maker, algorithms pre-select, curate, and personalize the decision environment. Yet these algorithms do not necessarily reduce uncertainty for the individual consumer with his or her information needs. Instead, non-transparent, dynamic, and responsive decision environments often seduce (dark patterns) or nudge towards certain options and can provide the individual with relatively inferior recommendations or choice sets as compared to static, non-responsive

F. G. Rebitschek (✉)

Harding Centre for Risk Literacy, Faculty of Health Sciences Brandenburg, University of Potsdam, Potsdam, Germany

 Max Planck Institute for Human Development, Berlin, Germany
e-mail: rebitschek@uni-potsdam.de

consumer decision environments with transparent options (Mathur, Mayer, & Kshirsagar, 2021). Data-driven behavioral control is unlikely to support informed consumer decision-making.

Informed decisions in Western industrialized countries (i.e., the German healthcare system by law (Deutscher Bundestag, 2013), constitute from an individual, who weighs the possible harms and benefits of alternative courses of action according to the best available evidence. Informed participation in algorithm-based environments (Gigerenzer, Rebitschek, & Wagner, 2018), moreover, requires continuous interaction with benefit-harm ratios that change dynamically due to external (e.g., the algorithm is modified by the provider) and internal (e.g., responding to one's past decisions) factors. Thus, besides understanding the benefits and harms of consuming or not consuming within a decision environment, grasping how the personal benefit-harm relationship changes dynamically can be crucial. Which strategies or rules need to be taught to consumers so they can reduce the uncertainty of challenging decision problems and are more likely making informed decisions?

Algorithms can support decision-making under uncertainty. One class of algorithms or models that boosts the decision-maker's competencies (Hertwig & Grüne-Yanoff, 2017) are fast-and-frugal decision trees (FFTs) (Martignon, Vitouch, Takezawa, & Forster, 2003). This type of algorithm aims to reduce a decision process to a handful of the most predictive combinations of features, termed cues. Consumers can robustly classify decision options (e.g., determine whether an informed decision is possible) by independently checking the presence or absence or level of those cues. Accordingly, the tree comprises classifications, decisions, or actions. Each cue comes with a branch either to the next cue or to an exit (e.g., a decision). In contrast to decision trees generally, FFTs involve no branching—apart from the last cue, which branches into two options (Martignon, Katsikopoulos, & Woike, 2008). From their structure, users glean that they can ignore further information, which makes FFTs a type of formal heuristics (Gigerenzer & Gaissmaier, 2011). Researchers in finance (Aikman et al., 2014), medicine (Green & Mehr, 1997), psychiatry (Jenny, Pachur, Williams, Becker, & Margrafm, 2013), and the military (Keller, Czienskowski, & Feufel, 2014) have shown that FFTs enable fast and reliable decisions—they perform similarly to more complex models (e.g., a logistic regression, random forest tree, and support vector machine).

As interpretable models that are transparent and educate those who use them, fast-and-frugal trees boost citizen empowerment (Harding Center for Risk Literacy, 2020a). They can be used as a graphically developed tree structure both digitally in apps, on websites, and analogue on posters and brochures. This facilitates them to be integrated in consumer decision-making. In a nutshell, fast-and-frugal trees lend the expert's view on a problem of uncertainty, providing a heuristic highly valid cue combination with which consumers separate the wheat from the chaff.

In the following section, I describe selected expert-driven decision-tree developments from the consumer research project RisikoAtlas (Harding Center for Risk Literacy, 2020a). The developed tools boost consumers when facing decisions under uncertainty across different domains: distinguishing between opinion and news; examining digital investment information; examining health information;

recognizing quality in investment advice, fake online stores, and unfair loan advice; detecting conflicts of interest in investment advice; controlling app data; enabling informed participation in bonus programs and credit scoring; informing telematics rate selection; and protecting data from employers and against personalized prices.

Methodology

As for any decision-support model development with instance-based learning, the developer must sample problem instances (cases), select decision cues (features) (and if faced with continuous cues, choose decision thresholds), and ensure that validation criteria are available. For FFT development specifically, the rank of cues and related exits is crucial and must be determined before validation. There are both manual (Martignon et al., 2008) and more complex construction algorithms using machine-learning methods (Phillips, Neth, Woike, & Gaissmaier, 2017) for developing FFTs.

However, the direct application of FFT-construction algorithms presupposes that a data set with problem instances, decision cues, and validation variables is available. Yet in consumer decision-making, usable data sets are exceptions in highly dynamic and algorithm-controlled decision environments. Accordingly, developers must sample problem instances from the environment, select decision cues based on experts and literature, and collect or investigate validation variables (cf. Keller et al., 2014). Here, I outline one expert-based development process (see Fig. 4.1).

Step 1: A developer must begin by identifying the problem for which he or she is designing decision support. Essential is the developer's reflection about what the decision tree should signal to the user. On the one hand, he or she must sharply define the classification target (the label, given a criterion). On the other hand, classifications and labelling can be abstract, with an open range of consequences, but also effective when connecting certain actions to the tree's exits. How strong should the recommendation be? How much of a norm should be conveyed by a tool that can make errors? Finally, it is essential to include at least one classification label with pre-defined staging.

Step 2: With regard to the later investigation of cue-target relationships, developers strive for an ecologically valid understanding of the prevalence of the potential cue structures by collecting material from a representative decision-making environment, such as real direct-2-consumer investment options, actual news pages, or descriptions or videos of real advisory situations. To sample sufficient material,

Fig. 4.1 Development pipeline according to a "case validity" FFT construction method. Source: Adapted from https://www.risikoatlas.de/en/consumer-topics/health/examining-health-information. Copyright 2020 by the Harding Center for Risk Literacy. Adapted with permission

this FFT development method ("case-based cue validity") requires the developer to consider not only the number of potential cues they must take into account with their modelling, but also the prevalence of the target, for example, what the decision tree should help its user to recognize.

Step 3: Having collected a basis of cases, developers must review potential cues that are predictive (or diagnostic) of the target. Experts, for instance in workshops, colleagues, laypersons, and literature such as scientific publications, white papers, authority reports, and experience reports help them collect a large set of candidate cues. It may also be useful to include new cues based on naïve theories or intuitions. A list of potential cues should then be examined in detail.

Each candidate cue must be understandable and verifiable under realistic conditions (time, costs, expertise) for potential users of the decision support tree. Otherwise, developers' risk ending up with cues that are highly predictive but inapplicable because they cannot be checked. Each additional cue requires more cases to enable robust development. As a rule of thumb, one can generally expect 20 to 50 cases for estimating the main effect of a single cue-target relationship, but considering cue combinations requires many more cases. If grouping similar cues, one has to consider that different levels of related cues in an underlying hierarchy could be predictive.

Step 4: To make use of the collected cases and cues, developers must code each cue status (presence, absence, level) for each of the cases. They can collect cue profiles with human resources or feature-detecting algorithms. Examining a hand-made process, one can identify many insights about the cues' actual usability for potential users. For instance, some cues may be too difficult for consumers to check and can be omitted from further development. As coding the status of each candidate cue across each case is very complex, cutting the number of cues early promises to be very efficient. However, before making a selection on the basis of how well consumers understand the cues, a statistical approach is usually cheaper and simpler.

Step 5: With respect to statistical analyses, developers must elicit whether or not a target criterion is met for each case (this label is required). This poses substantial challenges. An ideal approach would be testing of individual cases, so as to determine whether each case would satisfy a criterion. Very often, such effort is not feasible. For instance, it would imply examining hundreds of cases, each experimentally. As fast-frugal-trees lend the "expert's view," the wisdom of a small crowds of experts (Goldstein, McAfee, & Suri, 2014) can provide the best available evidence under realistic conditions (and the absence of data sets). The experts are provided with the cases, but they receive neither cue lists nor cue profiles. Several independent experts assess each case with respect to the criterion (they create the label). The median of a couple of their judgements proves to be robust when combining individual assessments (Galton, 1907).

Step 6: Based on the generated expert-based labels, statistical cue selection is possible. Simple statistical feature selection is already worthwhile after 100 cases, given limited cue numbers. Various tools are available for this, for example, the

boruta (Kursa & Rudnicki, 2010) and the caret package in R (Kuhn, 2008). With boruta, the developer can check an individual cue's validity in so-called random forests. If a cue behaves like a random number in a tree-based prediction of the label, he or she should not select it. Because the process is based on random sampling, the assessor should not ignore prior knowledge: If a known causal relationship exists, he or she should select the cue regardless and test it in repetition with more cases. The process of coding, scoring, and statistical feature selection can be done iteratively to more efficiently achieve a manageable set of robust cues. Cue coding effort and expert assessments depend on this set.

Step 7: If statistical cue selection no longer changes, given increasing case sets, it is worth modelling the decision tree based on the case profiles. To prevent overfitting, developers must separate predictive validity from description. They must separate a subset of cases, for example, two-thirds, from testing data, as training data for validation. FFT development and cross-validation can be done manually (e.g., ranking individual cues according to their predictive accuracy, posterior probability of being a true positive among all positive signals or altering positive and negative posteriors (Martignon et al., 2008), or with the help of machine learning algorithms (Phillips et al., 2017).

Step 8: Like any other classifier, FFTs perform a certain way in classification; they are more or less accurate or *efficacious*. They miss some cases in the real world and give a false alarm on others. To quantify their performance, developers can apply statistical cross-validation—that is, apply the decision tree randomly to repeated cases that form part of the testing dataset. An out-of-sample validation (external dataset) would be ideal. What quality is sufficient depends very much on the nature of the errors and the costs associated with the error.

Step 9: Finally, developers must test FFTs with regards to *effectiveness*. Randomized controlled experiments with the planned users are conducted for that purpose. Within experiments, one can for instance compare the decision-making of consumers who are presented with the decision tree with that of consumers who receive nothing or a standard information sheet.

Use Cases

Selecting Digital Health Information

Starting Point A comprehensive amount of health information on the web gives consumers the opportunity to learn about symptoms, benefits, or harms of medical interventions. Yet the quality of digital health information varies dramatically (Rebitschek & Gigerenzer, 2020). Misleading information leads to misperception of risks and prevents informed decisions (Stacey et al., 2017). Many sites have undeclared conflicts of interest. However, algorithmic curation of search results on both the web and news channels across social media platforms rarely comes with quality-dependent weighing (with

countermeasures having been implemented after this study). To prevent serious conse-
quences, consumers should be empowered to better recognize the quality of health
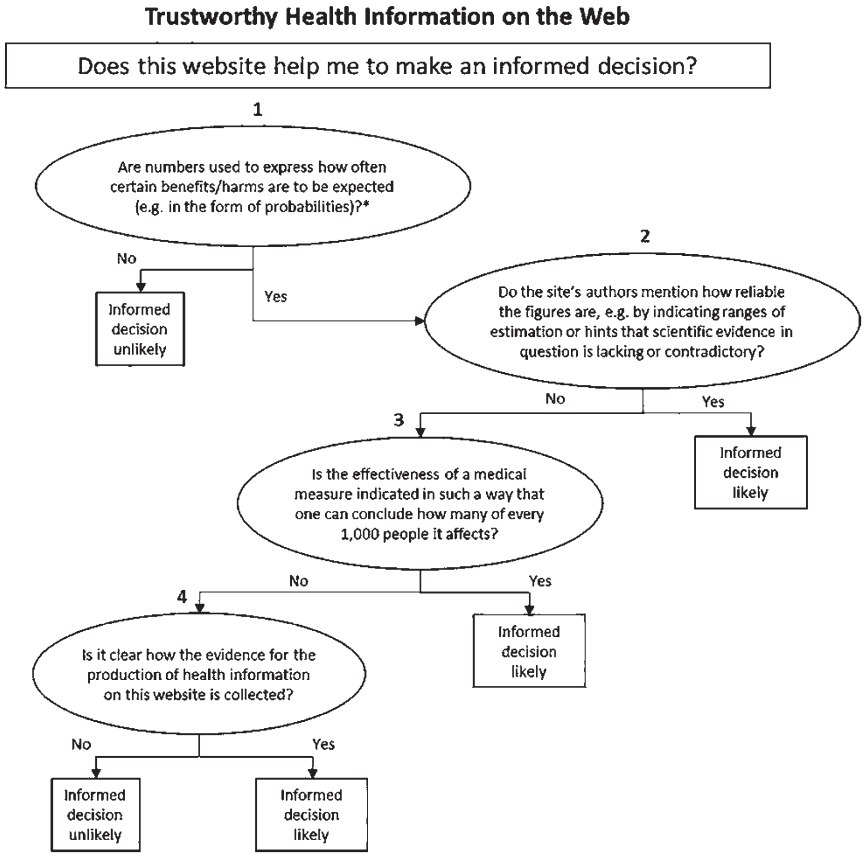information on the internet (Schaeffer, Berens, & Vogt, 2017).

Goal   How can one enable readers to distinguish between digital health information
that promotes informed decision-making and information that does not when they
do not even seek for the potential benefits and harms of decision options?

Cases, Cues, and Criteria   My team and I analyzed 662 pieces of health informa-
tion on German-language websites (Rebitschek & Gigerenzer, 2020). Of these, 487
were collected openly by experts, from Similarweb's health catalogue, and from
Google and Bing using medical condition terms (cf., (Hambrock, 2018) of diseases
and instrumental terms such as "How do I recognise X?". Another 175 pieces were
sampled by laypersons on given topics (vaccination against mumps, measles and
rubella; antibiotics for upper respiratory tract infections; ovarian cancer screening).
We artificially enriched the sample with randomly drawn pages from websites that
claim to follow the medical guideline for evidence-based health information in
Germany, which is an intentional oversampling compared to a random selection. We
aimed to predict the median classification judgments (label) of three experts per
piece about whether a piece enables or prevents informed health decision-
making (criterion). Experts stemmed from health information research, health
insurance companies, the Evidence-Based Medicine Network, and representatives
of health associations with professional experience in the field of health informa-
tion. We gave these experts no information about potential cues used in the study.

Development   By adhering to the evidence-based "Good Practice Health Information"
(EBM-Netzwerk, 2016) and the DISCERN standards, we identified 31 and 39 cues,
respectively, as verifiable by consumers. Elimination of redundant cues resulted in 65
cues. We conducted our cue selection stepwise using statistical methods, lay- and
expert comprehensibility, and usability. Finally, we considered 10 cues for modelling
with R. The final consumer tree with four cues is shown in Figure 4.2.

Interpretation   A warning means that one is probably unable to make an informed
decision based on the piece of health information in question. There can be many
reasons for this: It may be because essential information is being withheld. It may
be advertising or unprofessional design. In addition, following the decision tree may
lead one to the wrong conclusion, because the classifier is not perfect.

Validation of Efficacy   By cross-validating our health information set, we showed
its reliability. A cross validation of the identified decision tree resulted in a balanced
accuracy of 0.74. Following the FFT, users were warned in nine out of the ten health
information instances of which experts also stated they would not have been able to
reach an informed decision. Noteworthy, the decision tree only enabled users to
recognize six out of the ten cases of which experts stated they could have reached a
decision.

Trustworthy Health Information on the Web

Does this website help me to make an informed decision?

1

Are numbers used to express how often certain benefits/harms are to be expected (e.g. in the form of probabilities)?*

No

Informed decision unlikely

Yes

2

Do the site's authors mention how reliable the figures are, e.g. by indicating ranges of estimation or hints that scientific evidence in question is lacking or contradictory?

No

Yes

Informed decision likely

3

Is the effectiveness of a medical measure indicated in such a way that one can conclude how many of every 1,000 people it affects?

No

Yes

Informed decision likely

4

Is it clear how the evidence for the production of health information on this website is collected?

No

Yes

Informed decision unlikely

Informed decision likely

* In rare cases it is explained why there are no numbers available, which qualifies for the second question

Fig. 4.2 Fast-frugal tree to promote consumers' search for evidence-based information that supports informed health decisions. Source: Adapted from https://www.risikoatlas.de/en/consumer-topics/health/examining-health-information. Copyright 2020 by the Harding Center for Risk Literacy. Adapted with permission

Validation of Effectiveness With a lab-experimental evaluation (N = 204, 62% female, average age 40 years), we showed that the fast-and-frugal tree supports the assessment of health information. Independent experts assessed the lay people's findings in free internet searches for evidence-based health information. They rated users' search results on a four-point scale as worse in cases without a decision tree (2.7; a rather uninformed choice) than in those with one (2.4; a rather informed choice).

Selecting Digital Investment Options

Starting Point Consumers today commonly invest money on the internet, including into products of the so-called grey capital market. Direct-2-consumer investment options particularly lack the presence or advice of an expert. Potential investors must judge opportunities by relying on information either given or laterally (e.g., review pages). Many product providers are subject to less supervision than banks, for example, algorithmic advice, which aims to not be an advisor (according to German law) but is often labelled as a "robo-advisor." Transparency, even on the level required by law, is often absent, because the algorithms' architects intentionally hamper the weighing of potential gains and losses, and further risks.

Goal How can one enable potential investors to distinguish between digital investment options that are trustworthy, because they inform decision-making, and others aimed at blocking information, preventing the weighing of potential benefits and risks?

Cases, Cues, and Criteria My team and I analyzed 693 investment options on the web that were available to consumers in Germany. We searched for individual terms on Google and Facebook (bond, retirement provision, fund, investment, capital investment, return, savings, call money, securities), and after 100 options combined them with terms like interest, share, guarantee, gold, green, interest, precious metal, and ETF. We identified a further 180 cases through lay research. Furthermore, we manually sampled individual information on project offers on crowdfunding platforms. We did not include overview pages of individual banks on various capital investments (i.e., tabular listing of key figures on specific investment opportunities), advisory offers by banks or independent brokers, insurance companies, and financial managers. We aimed to predict the median classification judgments (label) of three experts per offer, whether an offer enables or prevents informed investing (criterion). 42 experts with academic or practical professional experience in the design of finance information evaluated the cases. We gave the experts no information about potential cues used in the study.

Development Based on various sources, we selected 138 cues, of which we considered 72 assessable in principle by laypersons after eliminating redundancies following an initial test. We conducted our cue selection stepwise using statistical methods, lay- and expert comprehensibility, and usability. Finally, we considered seven cues for modelling. The final consumer tree with four cues is shown in Figure 4.3.

Interpretation A warning means that informed investing is unlikely based on the offer in question. There can be many reasons for this: The provider could be interested in customers not making an informed decision, or the offer could be simply unprofessional. Also, following the decision tree can lead to a wrong conclusion,

Trustworthy Investment Opportunities on the Web

Can I make an informed investment at all here?
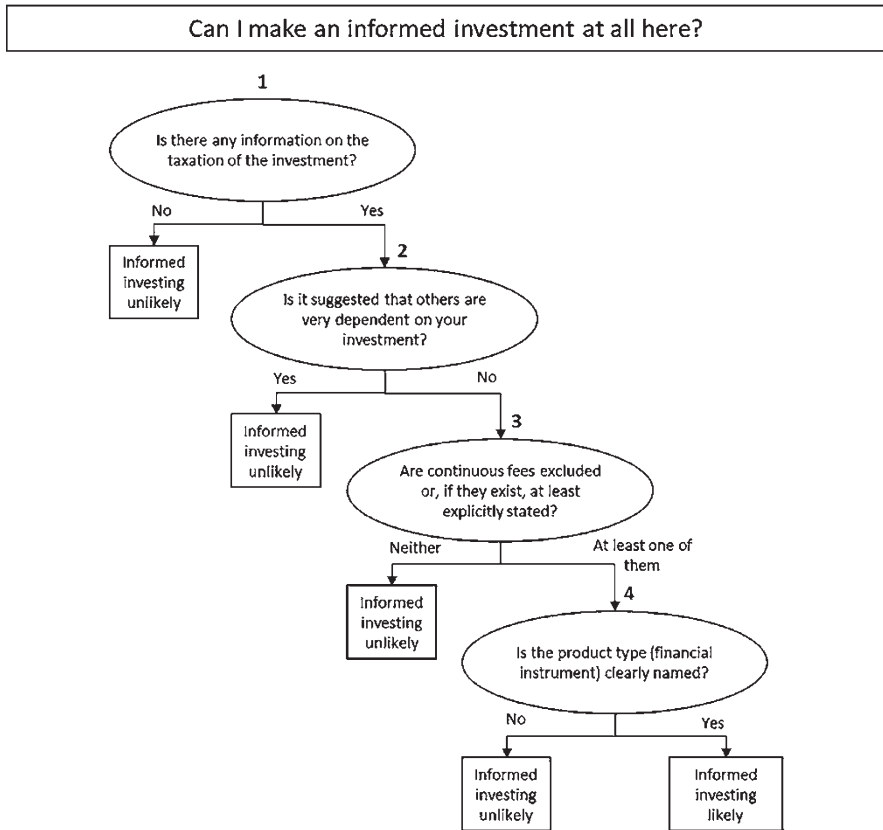
Fig. 4.3 Fast-frugal tree to promote consumers' search for trustworthy investment opportunities that promote informed investing. Source: Adapted from https://www.risikoatlas.de/en/consumer-topics/finance/examining-digital-investment-information. Copyright 2020 by the Harding Center for Risk Literacy. Adapted with permission

because the classifier is not perfect. Using the tree produces no insight into the quality of the offers themselves.

Validation of Efficacy By cross-validating the identified decision tree, we revealed a balanced accuracy of 0.78. Users are able to detect eight out of ten offers that enable informed investing, and reject seven out of ten because they do not enable informed investing.

Validation of Effectiveness With a lab-experimental evaluation (N = 204, 62% female, average age 40 years), we showed that an early version of the fast-and-frugal tree supports the search for consumer investment options on the web. Independent experts on finance investments assessed the lay people's findings of

investment options. They revealed that 385 out of 490 offers did not allow for informed investing. Although providing the tree did not let participants more often choose the rare options where they could invest on an informed basis, they at least became much more careful with investing in general, reducing the median initial hypothetical investments from 1000 to 500 € (retirement saving) and from 2500 to 1000 € (wealth accumulation).

Distinguishing News and Opinion Formats

Starting Point Social media users are more likely to like and share fake than real news, which is directly linked to the configuration of the algorithmic distribution. Consequently, algorithmic-based news coverage leads to misconceptions and makes social exchange more difficult. As fake news detection is challenging, a first step is to support users in distinguishing between news formats and opinion texts.

Goal How can one enable users to distinguish opinion formats and real news on social media and on websites?

Cases, Cues, and Criteria We fully analyzed 558 texts from German-language websites. Our topic selection based on fact checkers included "migration background," "chemtrails," "contrails," "Islam," "Muslims," "Israel," "cancer," "unemployed," "gender," "Russia," "VW," "left-wing extremism," "autonomists," "right-wing extremism," "money," and "climate." We complemented searches on Bing News, Google News, Facebook, Twitter, and those conducted with Google's "auto-complete" function with individual texts from the fake news portals described earlier. We were aiming to predict the median classification judgments (label) per text of three journalists with professional experience in print and digital media about whether the text's authors had satisfied or violated professional standards of the journalistic format "news" (criterion). We gave these experts no information about the potential cues used in the study.

Development Based on various sources, we collected 86 cues, of which we considered 50 to be basically verifiable by laypersons. We conducted our cue selection stepwise using statistical methods, expert comprehensibility, and usability. Finally, we used ten cues to model the satisfaction of journalistic standards. The final tree with four cues is shown in Figure 4.4.

Interpretation A warning means that the text violates professional journalistic standards of the news format. Examples are advertising, unprofessional texts, opinions such a commentary format, a satirical format, or so-called fake news. In some cases, those following the decision tree may reach the wrong conclusion, because the classifier is not perfect.

News or Opinion?

Does a text fail to meet journalistic requirements for news?
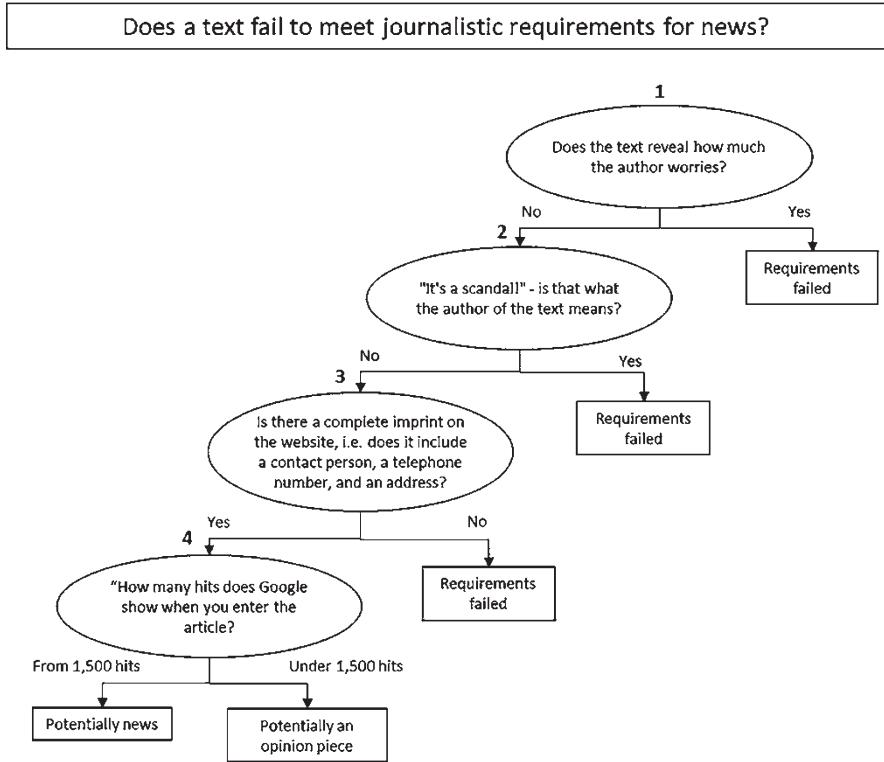
Fig. 4.4 Fast-frugal tree to help consumers classify news and opinion pieces. Source: Adapted from https://www.risikoatlas.de/en/consumer-topics/digital-world/distinguishing-between-opinion-and-news. Copyright 2020 by the Harding Center for Risk Literacy. Adapted with permission

Validation of Efficacy Cross-validating the decision tree, we reached a balanced accuracy of 0.76. Those following the decision tree recognized nine out of ten texts that were definitely not news as such, and similarly confirmed more than six out of ten real news texts.

Validation of Effectiveness With a lab-experimental evaluation (N = 204, 62% female, average age 40 years), we showed that 85% of laypeople applying the fast-and-frugal tree on 20 texts memorized all of the tree's cues with a short delay. Providing participants with the tree increased the overall classification accuracy from 74% to 78%, with a major advantage in confirming real news from 74% to 83%.

Discussion

Highly uncertain, non-transparent algorithm-controlled decision environments pose a threat to informed decision-making. Researchers have established that consumers are aware that the algorithms informing their decisions are imperfect, for example, in credit scoring, person analysis, and health behavior assessment (Rebitschek, Gigerenzer, & Wagner, 2021b). Yet consumers need more than awareness—they need applicable and educative tools (empowerment) to help reduce uncertainty.

With the help of three use cases, I have shown that fast-and-frugal decision trees can help users to distinguish quality-assured information from other pieces. Although efficacy in terms of absolute classification accuracies seems to be moderate, three arguments are relevant for their interpretation. First, to the best of my knowledge consumer support tools, at least in Germany, have never been validated with such empirical tests (e.g., for an overview over health information search support, see (Rebitschek & Gigerenzer, 2020). Thus, no one knows whether tools that are more accurate are even available. Second, a benchmark of absolute numbers is less relevant than a relative improvement over the current situation. This leads to the most important point, the validation in terms of effectiveness: The decision-makers in our studies made somewhat better choices and learned something given the moderate efficacy.

Thus, researchers within the field of consumer education should consider public engagement when developing uncertainty-reducing decision-support tools. FTTs are promising tools for boosting consumer competencies (Center for Adaptive Rationality, 2022), for instance for direct investment on the internet, in financial advice, or in the informed choice of a telematics tariff. They have been disseminated with a consumer app (Harding Center for Risk Literacy, 2020b). The next step has to be a pipeline for organizations that aim to protect citizens or consumers to develop and update them on a regular basis.

Even competence-promoting decision trees are always a temporary solution: Environments are dynamic and cues lose their predictive validity over time. Furthermore, transparent decision-support tools can be subject to gaming when information providers and decision architects consider merely fulfilling a desired cue status rather than actually improving the offers. Architects should not only consider including only causally related cues that cannot be gamed easily, but also subject their products to continuous updates.

As for any decision-support algorithm, the FFTs' limitations lie in their imperfect performance (classification errors). Therefore, actors must determine their follow-up actions carefully. In addition, the procedural fairness of information or products can be insufficient (i.e., when female consultants have a higher risk of misclassification), which needs to be controlled for every tree. Finally, decision-supporting tools, particularly algorithm-based decision making, set new norms. They inhere certain normativity. The importance of chosen criteria and cues can generalize, including to human decision-making. In addition, those introducing an algorithm cannot guarantee its implementative effectiveness in terms of side effects,

adverse events, and compensating behavior. Therefore, the most crucial factor for consumer empowerment algorithms—those which pre-select, curate, and personalize content, information, and offers—is regulatory examination. Empowerment and transparency have clear-cut limits, particularly in markets of data-driven behavioral prediction and control (i.e., consumer scoring (Rebitschek et al., 2021a), which helps define regulatory initiatives. This in turn emphasizes that regulation of knowledge and technology settles on the extent to which consumers become literate, to shape the participatory political and societal discourse on algorithm-based decision-making—the actual goal of empowerment.

References

Aikman, D., Galesic, M., Gigerenzer, G., Kapadia, S., Katsikopoulos, K., Kothiyal, A., Murphy, E., & Neumann, T. (2014). Taking uncertainty seriously: Simplicity versus complexity in financial regulation (Bank of England Financial Stability Paper No. 28). Retrieved from https://www.bankofengland.co.uk/-/media/boe/files/financial-stability-paper/2014/taking-uncertainty-seriously-simplicity-versus-complexity-in-financial-regulation.pdf?la=en&hash=2DE92C65BBF37630EE568ED475A4B2B2D996EE97

Center for Adaptive Rationality. (2022). Science of Boosting. Retrieved from https://scienceof-boosting.org/de/

Deutscher Bundestag. (2013. February 20). Gesetz zur Verbesserung der Rechte von Patientinnen und Patienten [Act to improve the rights of patients]. Berlin, Germany.

EBM-Netzwerk. (2016). *Gute Praxis Gesundheitsinformation: Ein Positionspapier des Deutschen Netzwerks Evidenzbasierte Medizin e.V.* [Good practice health information: Positionpaper by the German network for evidence-based medicine e.V.]. Berlin: Deutsches Netzwerk Evidenzbasierte Medizin e.V.. Retrieved from https://www.ebm-netzwerk.de/de/veroeffentlichungen/weitere-publikationen

Fritz, W., & Thiess, M. (1986). Das Informationsverhalten des Konsumenten und seine Konsequenzen für das Marketing [The consumer's information behavior and its consequences for marketing]. In F. Unger (Ed.), *Konsumentenpsychologie und Markenartikel* (pp. 141–176). Heidelberg: Physica-Verlag. https://doi.org/10.1007/978-3-642-93621-0_6

Galton, F. (1907). Vox Populi. *Nature, 75,* 450–451. https://doi.org/10.1038/075450a0

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology, 62,* 451–482. https://doi.org/10.1146/annurev-psych-120709-145346

Gigerenzer, G., Rebitschek, F. G., & Wagner, G. G. (2018). Eine vermessene Gesellschaft braucht Transparenz [A scored society needs transparency]. *Wirtschaftsdienst, 98,* 860–868. https://doi.org/10.1007/s10273-018-2378-4

Glückler, J., & Sánchez-Hernández, J. L. (2014). Information overload, navigation, and the geography of mediated markets. *Industrial and Corporate Change, 23,* 1201–1228. https://doi.org/10.1093/icc/dtt038

Goldstein, D. G., McAfee, R. P., & Suri, S. (2014). *The wisdom of smaller, smarter crowds.* Proceedings of the 15th ACM conference on economics and computation, 471–488. https://doi.org/10.1145/2600057.2602886

Green, L., & Mehr, D. R. (1997). What alters physicians' decisions to admit to the coronary care unit? *Journal of Family Practice, 45*(3), 219–227.

Hambrock, U. (2018). *Die Suche nach Gesundheitsinformationen: Patientenperspektiven und Marktüberblick* [The search for health information: patient perspectives and market overview]. Gütersloh: Bertelsmann Stiftung. https://doi.org/10.11586/2017053

Harding Center for Risk Literacy. (2020a). Project "RisikoAtlas" [Risk Atlas]. Retrieved from https://www.risikoatlas.de/en/risikoatlas-project/research

Harding Center for Risk Literacy. (2020b). *RisikoKompass* [Risk compass]. Retrieved from https://www.risikoatlas.de/en/tools-and-methods/informed-search-information/app-decision-support/risikokompass-risk-compass?view=methode

Hertwig, R., & Grüne-Yanoff, T. (2017). Nudging and boosting: Steering or empowering good decisions. *Perspectives on Psychological Science, 12,* 973–986. https://doi.org/10.1177/1745691617702496

Jenny, M. A., Pachur, T., Williams, S. L., Becker, E., & Margraf, J. (2013). Simple rules for detecting depression. *Journal of Applied Research in Memory and Cognition, 2*(3)*,* 149–157. https://doi.org/10.1037/h0101797

Keller, N., Czienskowski, U., & Feufel, M. A. (2014). Tying up loose ends: A method for constructing and evaluating decision aids that meet blunt and sharp-end goals. *Ergonomics, 57,* 1127–1139. https://doi.org/10.1080/00140139.2014.917204

Knight, F. H. (1921). *Risk, uncertainty, and profit.* Boston: Houghton Mifflin.

Kuhn, M. (2008). Building predictive models in R using the caret package. *Journal of Statistical Software, 28*(5), 1–26. https://doi.org/10.18637/jss.v028.i05

Kursa, M. B., & Rudnicki, W. R. (2010). Feature selection with the Boruta package. *Journal of Statistical Software, 36*(11), 1–13. https://doi.org/10.18637/jss.v036.i11

Lee, B.-K., & Lee, W.-N. (2004). The effect of information overload on consumer choice quality in an online environment. *Psychology & Marketing, 21*(3), 159–183. https://doi.org/10.1002/mar.20000

Martignon, L., Vitouch, O., Takezawa, M., & Forster, M. R. (2003). Naive and yet enlightened: From natural frequencies to fast and frugal decision trees. In D. Hardman & L. Macci (Eds.), *Thinking: Psychological perspective on reasoning, judgment and decision making* (pp. 189–211). Chichester: John Wiley & Sons. https://doi.org/10.1002/047001332X.ch10

Martignon, L., Katsikopoulos, K. V., & Woike, J. K. (2008). Categorization with limited resources: A family of simple heuristics. *Journal of Mathematical Psychology, 52*(6), 352–361. https://doi.org/10.1016/j.jmp.2008.04.003

Mathur, A., Mayer, J., & Kshirsagar, M. (2021). *What makes a dark pattern... dark? Design attributes, normative considerations, and measurement methods.* Proceedings of the 2021 CHI conference on human factors in computing systems, 1–18. https://doi.org/10.48550/arXiv.2101.04843

Phillips, N. D., Neth, H., Woike, J. K., & Gaissmaier, W. (2017). FFTrees: A toolbox to create, visualize, and evaluate fast-and-frugal decision trees. *Judgment and Decision Making, 12*(4), 344–368. https://doi.org/10.1017/S1930297500006239

Rebitschek, F. G., & Gigerenzer, G. (2020). Einschätzung der Qualität digitaler Gesundheitsangebote: Wie können informierte Entscheidungen gefördert werden? [Assessing the quality of digital health services: How can informed decisions be promoted?]. *Bundesgesundheitsblatt-Gesundheitsforschung-Gesundheitsschutz, 63,* 665–673. https://doi.org/10.1007/s00103-020-03146-3

Rebitschek, F. G., Gigerenzer, G., Keitel, A., Sommer, S., Groß, C., & Wagner, G. G. (2021a). Acceptance of criteria for health and driver scoring in the general public in Germany. *PLoS One, 16*(4), e0250224. https://doi.org/10.1371/journal.pone.0250224

Rebitschek, F. G., Gigerenzer, G., & Wagner, G. G. (2021b). People underestimate the errors made by algorithms for credit scoring and recidivism prediction but accept even fewer errors. *Scientific Reports, 11,* 20171. https://doi.org/10.1038/s41598-021-99802-y

Schaeffer, D., Berens, E.-M., & Vogt, D. (2017). Health literacy in the German population: Results of a representative survey. *Deutsches Ärzteblatt International, 114,* 53–60. https://doi.org/10.3238/arztebl.2017.0053

Stacey, D., Légaré, F., Lewis, K., Barry, M. J., Bennett, C. L., Eden, K. B., Holmes-Rovner, M., Llewellyn-Thomas, H., Lyddiatt, A., Thomson, R., & Trevena, L. (2017). Decision aids for people facing health treatment or screening decisions. *Cochrane Database of Systematic Reviews, 4,* CD001431. https://doi.org/10.1002/14651858.CD001431.pub5

Dr. Felix G. Rebitschek is Head Research Scientist and CEO of the Harding Center for Risk Literacy. Taking into account the nature of decision problems, his research aims at identifying accessible problem representations in risk communication in order to empower decision-makers, for example with educative interventions about digital health information. The associate researcher of the Max Planck Institute for Human Development in Berlin is cognitive psychologist with a focus on algorithm-based decision-making and he trains health professionals and managers on risk communication, risk literacy, and decision making under uncertainty internationally.