
What are you optimizing for?

Aligning Recommender Systems with Human Values

Jonathan Stray¹ Steven Adler¹ Dylan Hadfield-Menell^{2,3}

Abstract

We describe cases where real recommender systems were modified in the service of various human values such as diversity, fairness, well-being, time well spent, and factual accuracy. From this we identify the current practice of values engineering: the creation of classifiers from human-created data with value-based labels. This has worked in practice for a variety of issues, but problems are addressed one at a time, and users and other stakeholders have seldom been involved. Instead, we look to AI alignment work for approaches that could learn complex values directly from stakeholders, and identify four major directions: useful measures of alignment, participatory design and operation, interactive value learning, and informed deliberative judgments.

1. Introduction

Recommender systems are AI systems that present users with a tailored set of items based on factors such as past user behavior, user attributes, and features of the underlying items. They rank huge numbers of items in order to determine which content each user sees or interacts with on social media feeds, video platforms, new aggregators, and in online stores.

Classically, recommender designers have framed their goal as providing “relevant” items to each user. As these technologies have rolled out into society, it has become clear that the appropriate goals are often highly context-dependent, nuanced, and consequential. Previous work has argued for the incorporation of additional objectives including the promotion of accurate information (Stöcker, 2019), fairness (Barocas et al., 2019), and diversity of content (Castells

et al., 2015), and the reduction of addiction (Hasan et al., 2018)(Andreassen, 2015) and polarisation (Benkler et al., 2018). These represent values-driven decisions about what ‘we,’ the relevant stakeholders, do and don’t want recommended and how items should be ranked.

There are further, less-explored opportunities for values-driven design. A specification for a recommender system could include, e.g., that a news recommendation system to support values like informedness, autonomy, inclusiveness, participation, decentralization, representation, deliberation, and tolerance (Helberger, 2019) or that social media recommenders “provide positive incentives for encouraging more constructive conversation” (Wagner, 2019). Recommenders for books, music and video might want to promote cultural exploration and ensure that emerging artists gain exposure (Mehrotra et al., 2018). An online store may want its product recommendations to encourage the sale of products with a smaller carbon footprint (Rolnick et al., 2019) or otherwise support social and environmental goals. Recommenders have already been built around high-level goals like improving user well-being (Khwaja et al., 2019)(Stray, 2020) and even proposed as an aid to self-actualization (Knijnenburg et al., 2016).

In this paper, we propose that recommender system design can draw from the study of *the value alignment problem*: the problem of ensuring that an AI system’s behavior aligns with the values of the *principal* on whose behalf it acts (Hadfield-Menell & Hadfield, 2019). In the corresponding *recommender alignment problem*, the objective is to align recommendations with the goals of users, system designers, and society as a whole. In Section 2, we ground the recommender alignment problem with case studies of how the builders of large recommendation systems have responded to domain-specific challenges. We observe a common three phase approach to alignment: 1) relevant categories of content (e.g., clickbait) are *identified*; 2) these categories are *operationalized* as evolving labeled datasets; and 3) models trained off this data are used to adjust the system recommendations. In Section 3, we draw from AI alignment work to propose high-level approaches to the design of recommendation systems whose output embodies human values in a more nuanced, participatory and adapt-

¹Partnership on Artificial Intelligence, San Francisco, California, USA ²Department of Electrical Engineering and Computer Science, UC Berkeley, California, USA ³Center for Human-Compatible Artificial Intelligence, Berkeley, California, USA. Correspondence to: Dylan Hadfield-Menell <dhm@eecs.berkeley.edu>.

able way. We do not attempt to catalog the ethical issues surrounding recommenders (see e.g. (Milano et al., 2019a)) or to say how various values should be prioritized. Rather, we are interested in developing *affordances* that allow high-level, inclusive control of what a recommendation system recommends.

2. Case Studies

This section aims to ground our discussion of higher-level approaches to recommender alignment in concrete examples of diverse ways that large commercial recommender systems have been modified to be “better.”

2.1. Identifying unwanted content

Clickbait is one of the earliest recognized examples of recommender misalignment. Early recommender systems were often optimized to show the user the items they would be most likely to click on, which rewarded items with attention-grabbing titles but not necessarily quality content. In response, system designers identified alternate metrics that are harder to game and incorporated them into system training. For example, “dwell time”, the amount of time a user spends on a page before coming back to recommender results, has been used as ranking and personalization signal at many companies including Yahoo (Yi et al., 2014) and Facebook (Khalid El-Arini & Joyce Tang, 2014). Other simple metrics, such as the click-to-share ratio on Facebook posts (Khalid El-Arini & Joyce Tang, 2014), have also been used in practice to suppress clickbait.

However, hand-designed metrics are typically unable to capture the complexity of even relatively simple and uncontroversial concepts. Most large production systems today (including Flipboard (Cora, 2017) and Facebook (Peysakhovich & Hendrix, 2016)) instead train classifiers on human-labelled clickbait and use them to measure and decrease the prevalence of clickbait within the system. Recommenders that use these learned classifiers are more effective than those that use only hand-designed metrics. As we discuss below, we attribute this to the combination of a greater quantity of human feedback with a more flexible model.

There are a number of other types of content that one might not want to recommend broadly, such as hateful or harassing posts, sexual content, posts that advocate violence, or misinformation. These are too complex to detect with simple metrics derived from user behavior. Their precise definition is politically contested, and the decision to remove or down-rank this content must balance potential harm, local laws, and freedom of expression in complex, context-sensitive ways (York & Zuckerman, 2019). Hateful or harassing posts are challenging to identify by machine because of

the nuances of language and context but sometimes significant is available, for example if a comment occurs within a larger online discussion. A “toxicity” model created by Jigsaw, trained on human-labelled examples, has been tested at Twitter as part of a “healthy conversations” metric (Wagner, 2019), and used by the New York Times for comment triage prior to human review (Etim, 2017). However, such classifiers can misunderstand the significance of utterances in race- and gender-biased ways (Salganik & Lee, 2020).

Mis- and disinformation is a much harder problem for automated detection because it may require open-ended research, e.g. verifying a video by identifying and speaking to a witness (Silverman, 2020). Nonetheless, major platforms including Facebook (Tessa Lyons, 2018) now employ trained models in various aspects of their counter-misinformation efforts. These typically operate at the level of classifying domains, users, or near-duplicates of previously debunked items (Cora, 2017) (Facebook, 2020). The final source of truthfulness is human ratings of individual articles, photos, and videos. At Facebook this information is provided both by a network of professional fact-checking organizations and by users who flag content as “false news” (Thompson, 2018).

2.2. Promoting desirable outcomes

More recently, recommender system builders have attempted to identify content that furthers particular values-based goals. Goals addressed in practice include healthy public conversations, exposure for little-known creators or sellers, and the well-being of social media users. Rather than just identifying more or less desirable items, these richer conceptions of recommender alignment envision the support of positive values. Recommender systems could support democratic values in a variety of ways (Helberger, 2019), and encourage well-being (Khwaja et al., 2019), personal self-actualization (Knijnenburg et al., 2016), and human flourishing (Rodriguez & Watkins, 2009).

Although the empirical evidence for recommender systems’ role in creating “filter bubbles” is weak (Bruns, 2019), most discussions of recommenders and the public sphere hold diversity of viewpoint as an essential value (Helberger et al., 2018). Such studies typically consider diversity with respect to marginalized communities or left-right politics, but there are many other kinds of content diversity that recommender systems could promote such as topic or cultural diversity, and various concepts of representative fairness (Gao & Shah, 2020).

Spotify is notable for elaborating on the fairness and diversity issues faced by a music platform. Their recommendation system is essentially a two-sided market, since artists must be matched with users in a way that satisfies both, or neither will stay on the platform. Recommender systems

(and human cultural markets in general) often suffer from “superstar economics” where a small number of stars control the majority of the attention (Mehrotra et al., 2018). The company has also discussed the representation of different demographics, and the value of musical exploration, leading to a domain-specific definition: a playlist is “fair” if it contains songs from artists at different levels of popularity (Lalmas-Roelleke, 2019). This is a concrete example of the management of a multi-stakeholder recommendation system, a central topic in recommender alignment which is starting to be more explored (Abdollahpour et al., 2019).

Especially when filtering social media and news content, recommenders are key mediators in public discussion. Twitter has attempted to create “healthy conversation” metrics with the goal to “provide positive incentives for encouraging more constructive conversation” (Wagner, 2019). One conception of these metrics includes shared attention, shared reality, variety and receptivity, and a prototype has been built with models trained on human labels (Gadde & Gasca, 2018). The Dutch FG Mediagroep modified their news recommender to increase the consumption of “timely and fresh” content, as well as overall coverage, defined as the proportion of articles that were recommended to any user (Lu et al., 2020)

In the context of a mobile health app, researchers at Telefonica have shown that personalizing recommendations based on the user’s personality type, e.g. recommending social activities to extroverts, could result in a meaningful improvement in well-being outcomes (Khwaja et al., 2019). Facebook modified several of its recommendation systems in 2017-2018 to include consideration of “meaningful social interactions,” an explicit proxy for well-being (Stray, 2020). Such an intervention has potentially large effects because there are so many Facebook users, but no outcome information has so far been published.

2.3. Characterizing the State-of-the-Art

These case studies illustrate a common pattern in current recommender efforts: identifying and operationalizing abstract concepts like “healthy conversation” or “clickbait”, then adjusting the recommendation algorithm to change the observed prevalence of these concepts.

Identification: This is the phase where system designers become aware of a negative outcome associated with the system and identify a concept associated with it. For example, identifying that users are getting drawn into misrepresented videos for short periods of time and developing a definition of ‘clickbait’.

Operationalization: A concrete procedure is developed to identify instances of the abstract concept in the recommender system. Although this can be done through manual

review, most systems rely on some form of machine learning in order to scale. These models are typically trained from an evolving dataset where labels are provided by human workers or observable user behavior (e.g., flagging spam, responses to user surveys).

Adjustment: In this step, system designers modify the recommender system in order to increase or decrease the prevalence of the desired concept, as measured by the classification process defined in step 2. This can be done in a number of ways, most commonly: 1) *At training:* Adding concept prevalence as an additional training objective in a multi-objective training setup, as in (Zhao et al., 2019); 2) *At launch:* Tracking concept prevalence as a top-line metric and preferentially launching changes that improve it; 3) *At serving:* Directly add or remove items from the recommendation pool or change their rankings based on their classification, possibly after manual review.

3. Higher-level approaches to recommender alignment

A major problem with current practice is that it is fundamentally reactive. Problems are addressed one at a time, often after they are widespread. This is most troublesome when addressing one problem causes a new and harder-to-quantify problem, e.g. moving from optimizing for clicks to optimizing for time-in-app may reduce clickbait but could increase addiction. Even when there’s no risk of creating new problems, this approach simply cannot scale to handle the interests of billions of people embedded in millions of communities. Conducting detailed user research, designing new metrics, creating appropriate data sets, and training classification models is an iterative process that takes months if not years.

Moreover, the bottleneck in the problem identification step implies that only the problems most visible to system designers get addressed. This privileges issues that affect system designers and their social circles. Even more troubling is that it privileges the system designers’ moral and ethical views. We wish to address the problem of encoding appropriate values into recommender systems in a more general, participatory, and scalable way.

3.1. Useful definitions and measures of alignment

There is no single definition of value that a recommendation system should adhere to. This is in part a philosophical problem, in part the reality of political pluralism, and in part because value depends tremendously on context (Gabriel, 2020). But there are still many values that are widely shared, and useful instruments to evaluate human outcomes.

Suites of metrics for “well-being” are used in public policy (O’Donnell et al.) and many of those most applicable to AI

systems have been collected into an IEEE standard for well-being assessment (Schiff et al., 2020). Metrics for other high-level values, such as agency, have been developed in fields such as development economics (Alkire, 2008). These assessments rely on a combination of self-reported subjective states (“I am tired”), self-reported objective states (“I went to work today”), and statistics (“The unemployment rate is 5.4%”) (Stiglitz et al., 2018).

Many other metrics have been developed to evaluate recommenders, but that leaves open the question of which measures are most useful for value alignment in each domain. Standard evaluation data sets or protocols for recommenders in a particular domain will greatly accelerate progress on recommender alignment.

3.2. Participatory recommenders

The operators of recommender systems have often worked in a reactive mode, responding to issues after they are identified, and in the way that makes the most sense to them. It would be better if users could be effectively involved in the initial design of the system, had more control over routine operation, and good ways to escalate new issues.

Most real recommender systems must trade off between users or consumers, producers or sellers, and the recommender operator (Abdollahpouri et al., 2019). Recommendations may also affect non-users in a variety of ways, such as product recommendations that have adverse environmental consequences or driving directions that create congestion in otherwise calm neighborhoods (Kulynych et al., 2020). There sorts of issues are related to the classic economic problems of externalities and public goods, and are explicitly considered in multi-stakeholder recommendation research (Milano et al., 2019b). Users also have preferences around what recommendations *other* users see, such as when a someone flags something as misinformation, or when people work together to up- and down-vote items (York & Zuckerman, 2019).

The challenge here is in figuring out who the stakeholders are, how to involve them, and how to resolve their competing demands (Baum, 2020). One participatory design effort created an online matching system for a volunteer food delivery service, attempting “to enable stakeholders to construct a computational model that represents their views and to have those models vote on their behalf to create algorithmic policy” (Lee et al., 2019). Volunteer drivers, donors, recipients, and non-profit employees collaboratively designed the trade-offs between themselves and between efficiency and equity. This was done by interactive elicitation and modeling of individual preferences, followed by a Borda voting rule to decide the choice for the overall system. The authors report that this new algorithm both improved distributive outcomes and was perceived as fairer.

3.3. Interactive value learning

The recommenders with the most troubling side effects often have passive, low-agency, low-bandwidth interfaces like infinitely scrolling content feeds or autoplay. In these systems, user traces are evaluated by a combination of hand-coded and learned metrics. It is often unclear, even to system designers, what ‘values’ these metrics represent, as optimization-driven behavior can be counter-intuitive (Krakovna, 2018).

Direct user control of how values are enacted would provide both agency and transparency. This implies that recommenders must learn values interactively. There is a general sentiment that offline recommender benchmarks are decreasingly meaningful (Rohde et al., 2018), which further suggests that static values training data sets are inadequate.

Rather than depending on surveys for user feedback, we could be thinking about new kinds of ongoing interaction protocols. For example, probabilistic models like Inverse Reward Design (Hadfield-Menell et al., 2017) can be used to regularize recommender metrics and infer consistent combinations of metrics that generalize intent (Ratner et al., 2018). Interfaces designed around efficient imitation learning (Brown et al., 2020) and active reward learning (Biyik et al., 2020) might be a path towards more accessible, nuanced and predictable control of recommendation behavior.

3.4. Design around informed, deliberative judgment

Converging arguments in AI alignment (Tarleton, 2010), political philosophy (Daniels, 2020), and bioethics (Molewijk et al., 2008) suggest that informed, deliberative, and perhaps retrospective evaluations are of a higher quality than immediate judgements. Many people report that they watch more TV than they retrospectively endorse (Frey et al., 2007) and a similar type of behavior happens with some social media users (Andreassen, 2015). The notion of “time well spent” gets at similar ideas.

This is closely related to the common problem of focusing on short vs. long term outcomes, which is an issue with the use of metrics in AI generally (Thomas & Uminsky, 2020). One solution is to solicit informed, retrospective feedback. For example, one could show the user a summary of their recommendations and usage patterns in the past month, ask them to spend some time reflecting on it, and then rate how happy they are with this or how well it matches their goals. Versions of this idea have seen some use in the tech industry, sometimes known as “the regret test” (Eyal, 2018).

4. Conclusion

Recommendations systems are a cornerstone of the internet economy. They allow people to effectively interact with the

gigantic amount of information found in modern internet platforms. As recommendation systems proliferate, it is important to consider the impact these systems have and the values that drive ranking decisions. A crucial step in this process is to understand how these values are currently designed and updated. As our case studies illustrate, this can be largely described with a three step process where 1) outcomes are identified and reified as concepts; 2) the concepts are operationalized with a hand-coded or learned metric; and 3) this metric is used to adjust recommendation behavior. Looking to the future, we are excited about the potential for recommendation systems to better align with human values through the incorporation of well-being metrics, participatory approaches to objective design, interactive value learning, and optimizing for informed and deliberative preferences.

References

- Abdollahpouri, H., Adomavicius, G., Burke, R., Guy, I., Jannach, D., Kamishima, T., Krasnodebski, J., and Pizzato, L. Beyond Personalization: Research Directions in Multistakeholder Recommendation. *arXiv preprint arXiv:1905.01986*, 2019.
- Alkire, S. Concepts and Measures of Agency. *Oxford Poverty & Human Development Initiative*, 2008. URL <https://ora.ox.ac.uk/objects/uuid:cdecba447c-43b7-8e3f-851517b5ff97>.
- Andreassen, C. S. Online Social Network Site Addiction: A Comprehensive Review. *Current Addiction Reports*, 2(2):175–184, 2015.
- Barocas, S., Hardt, M., and Narayanan, A. Fairness and machine learning, 2019. URL <https://fairmlbook.org/>.
- Baum, S. D. Social choice ethics in artificial intelligence. *AI and Society*, 35(1):165–176, 2020.
- Benkler, Y., Faris, R., and Roberts, H. *Network propaganda: Manipulation, disinformation, and radicalization in American politics*. Oxford University Press, 2018.
- Biyik, E., Huynh, N., Kochenderfer, M. J., and Sadigh, D. Active preference-based gaussian process regression for reward learning. *arXiv preprint arXiv:2005.02575*, 2020.
- Brown, D. S., Coleman, R., Srinivasan, R., and Niekum, S. Safe imitation learning via fast bayesian reward inference from preferences. *arXiv preprint arXiv:2002.09089*, 2020.
- Bruns, A. *Are Filter Bubbles Real?* Polity, 2019. ISBN 978-1509536443.
- Castells, P., Hurley, N. J., and Vargas, S. Novelty and diversity in recommender systems. In *Recommender systems handbook*, pp. 881–918. Springer, 2015.
- Cora, M. V. Detecting Trustworthy Domains. *Flipboard*, 2017. URL <https://engineering.flipboard.com/2017/04/domainranking>.
- Daniels, N. Reflective equilibrium. In Zalta, E. N. (ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2020 edition, 2020.
- Etim, B. The Times Sharply Increases Articles Open for Comments, Using Google’s Technology, June 2017. URL <https://www.nytimes.com/2017/06/13/insider/have-a-comment-leave-a-comment.html>.
- Eyal, N. Want to Design User Behavior? Pass the ‘Regret Test’ First, 2018. URL <https://www.psychologytoday.com/us/blog/automatic-you/201803/want-design-user-behavior-pass-the-regret-test-first>.
- Facebook. Using AI to detect COVID-19 misinformation and exploitative content, 2020. URL <https://ai.facebook.com/blog/using-ai-to-detect-covid-19-misinformation-and-exploitative-content>.
- Frey, B. S., Benesch, C., and Stutzer, A. Does Watching TV Make Us Happy? *Journal of Economic Psychology*, 28(3):283–313, 2007.
- Gabriel, I. Artificial intelligence, values and alignment. *arXiv preprint arXiv:2001.09768*, 2020.
- Gadde, V. and Gasca, D. Measuring healthy conversation. Twitter, 2018. URL https://blog.twitter.com/en_{_}us/topics/company/2018/measuring_{_}healthy_{_}conversation.html.
- Gao, R. and Shah, C. Toward creating a fairer ranking in search engine results. *Information Processing and Management*, 57(1), 2020.
- Hadfield-Menell, D. and Hadfield, G. K. Incomplete contracting and AI alignment. *AIES 2019 - Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 417–422, 2019.
- Hadfield-Menell, D., Milli, S., Abbeel, P., Russell, S. J., and Dragan, A. Inverse reward design. In *Advances in Neural Information Processing Systems 30*, pp. 6765–6774, 2017.

- Hasan, M. R., Jha, A. K., and Liu, Y. Excessive use of online video streaming services: Impact of recommender system use, psychological factors, and motives. *Computers in Human Behavior*, 80:220–228, 2018.
- Helberger, N. On the Democratic Role of News Recommenders. *Digital Journalism*, 7(8):993–1012, 2019.
- Helberger, N., Karppinen, K., and D’Acunto, L. Exposure diversity as a design principle for recommender systems. *Information Communication and Society*, 21(2):191–207, 2018.
- Khalid El-Arini and Joyce Tang. Click-baiting, 2014. URL <https://about.fb.com/news/2014/08/news-feed-fyi-click-baiting/>.
- Khwaja, M., Ferrer, M., Iglesias, J. O., Faisal, A. A., and Matic, A. Aligning daily activities with personality: towards a recommender system for improving wellbeing. In *Proceedings of the 13th ACM Conference on Recommender Systems*, pp. 368–372, 2019.
- Knijnenburg, B. P., Sivakumar, S., and Wilkinson, D. Recommender systems for self-actualization. *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys ’16)*, pp. 11–14, 2016.
- Krakovna, V. Specification gaming examples in ai, 2018. URL <https://vkrakovna.wordpress.com/2018/04/02/specification-gaming-examples-in-ai/>.
- Kulynych, B., Overdorf, R., Troncoso, C., and Gürses, S. Pots: Protective optimization technologies. *arXiv preprint arXiv:1806.02711v6*, 2020.
- Lalmas-Roelleke, M. Recommending and searching @ Spotify, 2019. URL <https://www.slideshare.net/mounialalmas/recommending-and-searching-spotify/21>.
- Lee, M. K., Kusbit, D., Kahng, A., Kim, J. T., Yuan, X., Chan, A., See, D., Noothigattu, R., Lee, S., Psomas, A., and Procaccia, A. D. Webuildai: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction (CSCW)*, 3, 2019.
- Lu, F., Dumitrache, A., and Graus, D. Beyond Optimizing for Clicks: Incorporating Editorial Values in News Recommendation. *arXiv preprint arXiv:2004.09980*, 2020.
- Mehrotra, R., McInerney, J., Bouchard, H., Lalmas, M., and Diaz, F. Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. *International Conference on Information and Knowledge Management, Proceedings*, pp. 2243–2252, 2018.
- Milano, S., Taddeo, M., and Floridi, L. Recommender Systems and their Ethical Challenges. 2019a. URL <https://ssrn.com/abstract=3378581>.
- Milano, S., Taddeo, M., and Floridi, L. Ethical aspects of multi-stakeholder recommendation systems. 2019b. URL <https://ssrn.com/abstract=3378581>.
- Molewijk, A. C., Abma, T., Stolper, M., and Widdershoven, G. Teaching ethics in the clinic. The theory and practice of moral case deliberation. *Journal of Medical Ethics*, 34(2):120–124, 2008.
- O’Donnell, G., Deaton, A., Durand, M., Halpern, D., and Layard, R. Wellbeing and Policy. *Legatum Institute*.
- Peysakhovich, A. and Hendrix, K. Further Reducing Click-bait in Feed. Facebook, 2016. URL <https://about.fb.com/news/2016/08/news-feed-fyi-further-reducing-clickbait-in-feed/>.
- Ratner, E., Hadfield-Mennell, D., and Dragan, A. D. Simplifying reward design through divide-and-conquer. In *Robotics: Science and Systems*, 2018.
- Rodriguez, M. A. and Watkins, J. H. Faith in the algorithm, part 2: Computational eudaemonics. In *International Conference on Knowledge-Based and Intelligent Information and Engineering Systems (KES)*, pp. 813–820, Berlin, 2009. Springer.
- Rohde, D., Bonner, S., Dunlop, T., Vasile, F., and Karatzoglou, A. RecoGym: A Reinforcement Learning Environment for the problem of Product Recommendation in Online Advertising. *arXiv preprint arXiv:1808.00720*, 2018.
- Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., Ross, A. S., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A., et al. Tackling climate change with machine learning. *arXiv preprint arXiv:1906.05433*, 2019.
- Salganik, M. J. and Lee, R. C. To Apply Machine Learning Responsibly, We Use It In Moderation. *New York Times Open*, 2020. URL <https://open.nytimes.com/to-apply-machine-learning-responsibly-we-use-it-in-moderation-d001f49e0644>.
- Schiff, D., Ayesh, A., Musikanski, L., and Havens, J. C. Ieee 7010: A new standard for assessing the well-being implications of artificial intelligence. *arXiv preprint arXiv:2005.06620*, 2020.
- Silverman, C. (ed.). *The Verification Handbook for Disinformation And Media Manipulation*. European Journalism Center, 2020. URL <https://datajournalism.com/read/handbook/verification-3>.

- Stiglitz, J. E., Fitoussi, J.-P., and Durand, M. *Beyond GDP*. OECD, 2018. URL <https://www.oecd-ilibrary.org/content/publication/9789264307292-en>.
- Stöcker, C. How Facebook and Google Accidentally Created a Perfect Ecosystem for Targeted Disinformation. In Grimme, C., Preuss, M., Takes, F. W., and Waldherr, A. (eds.), *Disinformation in Open Online Media*, pp. 129–149. Springer, 2019.
- Stray, J. Aligning AI to Human Values means Picking the Right Metrics. Partnership on AI, 2020. URL <https://medium.com/@PartnershipAI/aligning-ai-to-human-values-means-picking-the-right-metrics-855859e6f047>.
- Tarleton, N. Coherent Extrapolated Volition: A Meta-Level Approach to Machine Ethics. *Machine Intelligence Research Institute*, 2010.
- Tessa Lyons. Increasing Our Efforts to Fight False News. *Facebook*, 2018. URL <https://about.fb.com/news/2018/06/increasing-our-efforts-to-fight-false-news/>.
- Thomas, R. L. and Uminsky, D. Reliance on Metrics is a Fundamental Challenge for AI. In *Ethics of Data Science Conference*, 2020.
- Thompson, N. Facebook Opens Up About False News. *Wired*, 2018. URL <https://www.wired.com/story/exclusive-facebook-opens-up-about-false-news/>.
- Wagner, K. Inside Twitter’s ambitious plan to change the way we tweet. *Vox*, 2019. URL <https://www.vox.com/2019/3/8/18245536/exclusive-twitter-healthy-conversations-dunking-research-product-incentives>.
- Yi, X., Hong, L., Zhong, E., Liu, N. N., and Rajan, S. Beyond clicks: Dwell time for personalization. *Proceedings of the 8th ACM Conference on Recommender Systems (RecSys ’14)*, pp. 113–120, 2014.
- York, J. and Zuckerman, E. Moderating the Public Sphere. In Jørgensen, R. F. (ed.), *Human Rights in the Age of Platforms*, chapter 6. MIT Press, 2019.
- Zhao, Z., Chi, E., Hong, L., Wei, L., Chen, J., Nath, A., Andrews, S., Kumthekar, A., Sathiamoorthy, M., and Yi, X. Recommending what video to watch next: A Multitask Ranking System. In *Proceedings of the 13th ACM Conference on Recommender Systems (RecSys ’19)*, pp. 43–51, 2019.