

# First Do No Harm: Considering and Minimizing Harm in Recommender Systems Designed for Engendering Health

Jennifer D. Ekstrand<sup>2</sup> and Michael D. Ekstrand<sup>1,2</sup>

<sup>1</sup>Dept. of Computer Science  
Texas State University  
San Marcos, TX USA

<sup>2</sup>Dept. of Computer Science  
Boise State University  
Boise, ID USA

{jdekstrand,ekstrand}@acm.org

## ABSTRACT

In this paper, we consider some ways in which recommender systems designed to engender health may bring harm to their users, or others they impact, and discuss how the ethic of beneficence applies to health-related recommendations. We consider nutrition-oriented recommendation, examining some of the adverse outcomes that may arise as the recommender interacts with people in broad, real-world scenarios, and present a set of problems and some suggestions that researchers and practitioners may consider as they seek to maximize the positive impact of health recommendation.

## Keywords

Health; recommender systems; CSCW; participatory design; ethics; beneficence

## 1. INTRODUCTION

“First do no harm.” That utterance is one of the most widely-recognized statements of the fundamental bioethics principal of *non-maleficence*. Non-maleficence is one component of the principle of *beneficence* that is foundational to human subjects research [13] and is useful as a guiding light for many forms of intervention and technology.

We are pleased to see the growth in application of recommender systems to promoting healthy decision-making. There is great potential for encouraging better nutrition or, eventually, aiding doctors in diagnosis. However, we would like to promote consideration of those who could be harmed by the very systems designed to promote health. We submit that it is important to carefully consider—ideally with relevant experts—the various ways in which a technological health intervention could be misused or interact badly with its context of use, so that we can be justified in our assessment that its risks are minimal and outweighed by the benefits. Such an assessment – required by research ethics review in the U.S. – must be done with appropriate due diligence and imagination.

How might an unscrupulous employer use data from a fitness tracker—for example observing a change in activity due to pregnancy? What impact might health gamification have on someone with a tendency to addiction? How might someone coping with the frustrating experience of injury and recovery respond to default activity level suggestions? Will someone construe taking a recommendation to make a small healthy choice—such as taking stairs instead of the elevator—as an excuse to justify a less-healthy choice such as eating a disproportionate dessert? What economic harm might result from unfounded recommendations?

The direct users of health technology are not the only people it affects. Does a social wellness effort focused on steps alienate a person who relies on a wheelchair? Non-participation can directly lead to social exclusion: in one study, a non-participant was frustrated by being unable to join a campaign many of her co-workers were participating in, and that her exclusion meant she “couldn’t contribute or participate in that conversation” [5]. If someone

cheats on reporting in a social context, how might that influence their peers?

There has been significant work in food recommendation [4,6,10], so in this paper, we focus on the case of nutrition recommendation, examining various ways in which the recommender could result in adverse outcomes. These harms arise from a variety of sources, including misuse of the recommender, interaction with users whose health situation is outside the scope of the originally intended audience, and users whose living situations make it difficult to follow the recommendations as envisioned.

Within the context of nutritional recommendation, we attempt to identify some general categories of potential harm that we believe health recommendation researchers should consider and propose principles for mitigating these issues. We do not allege that current research never incorporates these principles; we are arguing that regular consideration of them in health recommendation research and design can help to produce systems that, to the best of our ability, do no harm.

## 2. CASE STUDY: NUTRITION

So far, many attempts toward using recommender systems for health have focused on nutrition—particularly in terms of reducing obesity or managing diabetes. However, defining “healthy” and “nutritious” in a world of varied contexts may result in inappropriate advice.

### 2.1 Non-Actionable Recommendations

A nutrition-focused recommender may be frustrating if **recommendations are not actionable**. If recommended recipes are too expensive, too time-consuming, too complicated, or contain ingredients that are not easily obtained, a parent may feel guilt for not providing the healthiest food for their family. Failing to meet complex needs may increase stress or anxiety rather than helping. Unusable recommendations may discourage users from trying to meet goals rather than helping them to succeed. *Participatory design* is one mechanism that can be used to give users the opportunity to influence a system designed for their needs [1,3,11].

Diabetes management has been one domain targeted for healthy food recommendation[9], and with good reason—all diabetics need to control the amount of carbohydrates they consume. However, not all variations of diabetes are the same. Diabetes mellitus type 2 is the most common form, and it is highly connected to obesity. However, a woman with temporary gestational diabetes needs to manage her blood sugar while consuming the appropriate nutrients for pregnancy rather than attempting to lose weight. For someone with cystic fibrosis related diabetes, maintaining enough weight is a challenge, and an appropriate diet should be high in calories, fats, and salt while counting carbs [14]. Recommender systems should *work with health professionals* who can speak to diverse people’s individual nutritional needs rather than attempting to replace them with generic standards of healthy food.

Well-designed recommender systems present an opportunity to help people with very specific dietary needs, but a **poorly-designed system may increase the stress of finding appropriate food**. Dealing with immune system disorders such as celiac disease and eosinophilic esophagitis requires significant effort to avoid foods which cause allergic reaction. A well designed system should *empower users to tailor the characteristics* of their recommendations.

Users may have **difficulty following recommendations**. Perhaps there is difficulty obtaining the recommended food from living in a food desert. Perhaps finances or time are limited. It may be that someone doesn't know how to prepare food in the manner specified by the recommendation. Ideally, a recommender system should *facilitate use of other resources*: a recipe recommendation might link to online tutorials or local cooking classes, or tag recipes that can make use of local food supplemental programs, such as the WIC program in the United States.

## 2.2 Encouragement of unhealthy behavior

It is possible that recommender systems may **inadvertently encourage unhealthy behavior**. Gamification may incentivize short-term results over sustainable long-term habits—promoting yo-yo dieting. Popular diets—such as the gluten-free trend—may be harmful when inappropriately used [8].

Nutrition is not a problem that **exists in isolation**. One way in which nutritional needs can be complicated is in interaction with mental health. Lifestyle interventions for those diagnosed with serious mental illness—who are also at increased risk of obesity—are particularly difficult. Yarborough et al. suggest “it is important that such interventions can be implemented in community-based settings” [12]. Recommender systems should *avoid undermining existing best practices*. Removing the social aspect of lifestyle interventions may be harmful.

Recent research has observed pro-anorexia communities online [2,7]. It is difficult to empower users to reach their own goals without **enabling destructive behavior** and, as one study mentioned, outright censorship may make intervention more difficult [2]. However, when designing recommender systems, we should attempt to predict how people with eating disorders may interact with them. Attempt to design in ways that encourage healthy behavior *beyond simplistic metrics* such as minimizing caloric intake or maximizing weight loss.

**Obesity may not be someone's most pressing health concern**. If weight-gain is a side effect of an important medicine (common in mental health) or an excuse to continue smoking, systems placing too much emphasis on weight may encourage unhealthy behavior to control weight. Where appropriate, *consider nutrition in the overall context of health*.

## 3. CONCLUSION

We advocate for systems to be designed in a process that includes imagination and empathy, includes intended users and considers their contexts, and partners with knowledgeable domain experts. Measuring overall well-being in long-term studies, while significantly more difficult than individual, isolated studies and small system trials, will also help to assess whether our innovations are truly helping the people they affect.

Recommender systems have great potential to help people live healthier lives, both through direct interaction and by empowering their health care professionals. To ensure that systems achieve this potential, it is necessary to carefully consider who might be harmed in order to minimize risk while maximizing benefits.

## 4. REFERENCES

- [1] Bisafar, F.I. and Parker, A.G. Confidence & Control: Examining Adolescent Preferences for Technologies That Promote Wellness. *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, ACM (2016), 160–171.
- [2] Chancellor, S., Pater, J.A., Clear, T., Gilbert, E., and De Choudhury, M. #Thyghgapp: Instagram Content Moderation and Lexical Variation in Pro-Eating Disorder Communities. *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, ACM (2016), 1201–1213.
- [3] Ekstrand, M.D. and Willemsen, M.C. Behaviorism is Not Enough: Better Recommendations through Listening to Users. *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*, ACM (2016), 4.
- [4] Freyne, J. and Berkovsky, S. Recommending Food: Reasoning on Recipes and Ingredients. In P. De Bra, A. Kobsa and D. Chin, eds., *User Modeling, Adaptation, and Personalization*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010, 381–386.
- [5] Gorm, N. and Shklovski, I. Steps, Choices and Moral Accounting: Observations from a Step-Counting Campaign in the Workplace. *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, ACM (2016), 148–159.
- [6] Norman, G.J., Zabinski, M.F., Adams, M.A., Rosenberg, D.E., Yaroch, A.L., and Atienza, A.A. A Review of eHealth Interventions for Physical Activity and Dietary Behavior Change. *American Journal of Preventive Medicine* 33, 4 (2007), 336–345.e16.
- [7] Pater, J.A., Haimson, O.L., Andalibi, N., and Mynatt, E.D. “Hunger Hurts but Starving Works”: Characterizing the Presentation of Eating Disorders Online. *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, ACM (2016), 1185–1200.
- [8] Reilly, N.R. The Gluten-Free Diet: Recognizing Fact, Fiction, and Fad. *The Journal of Pediatrics*, (2016).
- [9] Rokicki, M., Herder, E., and Demidova, E. What's On My Plate: Towards Recommending Recipe Variations for Diabetes Patients Markus Rokicki, Eelco Herder, Elena Demidova L3S Research Center, Hannover, Germany, 2015.
- [10] Said, A. and Bellogín, A. You are What You Eat! Tracking Health Through Recipe Interactions. *RSWeb@ RecSys*, (2014).
- [11] Schuler, D. and Namioka, A., eds. *Participatory Design: Principles and Practices*. CRC / Lawrence Erlbaum Associates, Hillsdale, N.J, 1993.
- [12] Yarborough, B.J.H., Leo, M.C., Stumbo, S., Perrin, N.A., and Green, C.A. STRIDE: a randomized trial of a lifestyle intervention to promote weight loss among individuals taking antipsychotic medications. *BMC Psychiatry* 13, (2013), 238.
- [13] *The Belmont Report*. U.S. Department of Health, Education, and Welfare, 1979.
- [14] Carol Brunzell, Dana S. Hardin, Anne Kogler, Antoinette Moran, and Terri Schindler. *Managing Cystic Fibrosis-Related Diabetes (CFRD): An Instruction Guide for Patients and Families*. Cystic Fibrosis Foundation, 2015.