We are often told to eat healthily. But how does the average person translate this into practise? We know that universally people can form healthiness judgments of individual foods without knowing anything about their nutritional composition. However, to date, perceptions of food healthiness have been predominately investigated in relation to how well they match up with government recommended nutrient guidelines. This ignores the vast information people are exposed to about foods from various media sources that shapes their beliefs; information that might even be at odds with a food's objective nutrient content. Recent advancements in computational linguistics mean we can capture this information (aka food associations) using large-scale natural language data. These machine-learning techniques enable us to develop generalizable models that can make judgment predictions for nearly any food item.

Our first and main model is our Vector Representation Model, which uses 300-dimensional vector representations to quantify these complex food associations. The second, our baseline Nutrient Model, predicts judgments using calorie content, amounts of nutrients (fat, saturates, sugar, salt, sodium and protein) per 100g, and traffic light colour coding criteria to form a 10-dimensional vector. Lastly, we concatenated the vectors from the Nutrient Model and Vector Representation Model in our Combined Model to form a 310-dimensional vector. We present a series of studies in which we assess the importance of nutrient and non-nutrient attributes of foods individually and together, with the goal of uncovering what associations underpin all food healthiness judgments. We also considered the effects of nutritional knowledge and different nutrient labelling formats on healthiness judgments by training our models on different sets of data. In total, our models were fitted to data from five studies (n = ~200 participants in each). All participants rated the healthiness of the same randomly selected foods on a scale ranging from -100 (most unhealthy) to +100 (most healthy). This allowed us to assess differences between members of the general public and nutritional experts (dietitians) shown food names only; as well as differences in judgments between individuals exposed to calorie information, monochrome nutrient labelling and traffic light labelling.

Using regularized linear functions, we attempted to predict the actual (aggregate) healthiness ratings of out-of-sample foods. Our findings showed that the vector representation model performed very well across all studies and conditions, with predictive accuracy (r^2) ranging from 0.65 to 0.77. By comparison, the predictive accuracy of the nutrient model was lower and much more variable (ranging from 0.33 to 0.77). Only in the case of the most complex label, the Traffic Light label (Study 4), did the nutrient model perform better than the vector representation model. We also found that the combined model performed better than the individual models in 7 out of 8 cases. The highest accuracy was achieved in the case of Traffic Light labels, with r^2 of 0.91, suggesting associations capture something unique about the way people make healthiness judgments. Together, these results show that associations with food names play an important role in people's judgments of food healthiness, often more so than the nutritional composition of the food. Whilst the dimensions of our word vector model are not directly interpretable, we were able to identify words that relate to the concept of food healthiness in semantic space. This was achieved by passing the vector representations of common words through our model trained on participants' food healthiness judgments. Our model predictions found that high ratings of healthiness correspond to words related to agriculture and nature (e.g. "crop", "harvest", and "agricultural"). These associations continue to exert a strong effect on people's judgements even when nutritional labelling is provided. As such, these results show a new area where interventions could be targeted to improve people's judgments and choices of healthy foods.

Our data driven and bottom up approach provides deeper insights into healthiness judgment formation than has been possible in a traditional study. This study demonstrates the foundations of an easily accessible tool that can be built upon and extended to help registered dietitians and policy makers alike evaluate current and future health interventions.