



Can we Use Conceptual Spaces to Model Moral Principles?

Steven Verheyen^{1,2}  · Martin Peterson³

Published online: 14 July 2020
© The Author(s) 2020

Abstract

Can the theory of conceptual spaces developed by Peter Gärdenfors (2000, 2014) and others be applied to moral issues? Martin Peterson (2017) argues that several moral principles can be construed as regions in a shared similarity space, but Kristin Shrader-Frechette (2017) and Gert-Jan Lokhorst (2018) question Peterson's claim. They argue that the moral similarity judgments used to construct the space are underspecified and subjective. In this paper, we present new data indicating that moral principles can indeed be construed as regions in a multidimensional conceptual space on the basis of moral similarity judgments. Four hundred and seventy-five students taking a course in engineering ethics completed a survey in which they were presented with ten cases (moral choice situations) featuring ethical issues related to technology and engineering. Participants were asked to judge the moral similarity of each pair of cases (45 comparisons) and to select which moral principle (from a list of five alternatives plus a sixth option: "none of the principle listed here") they believed should be applied for resolving the case. We used interval multidimensional scaling (MDS) as well as individual differences scaling (INDSCAL) for analyzing the moral similarity judgments. Despite noteworthy individual variations in the judgments, the five moral principles included in the study were discernable in the aggregate multidimensional spaces, even for participants with no previous exposure to the principles. Participants tended to apply the same moral principles to cases rated as morally similar. Our overall conclusion is that moral similarity judgments, and their representation in multidimensional spaces, can help us identify moral principles that are relevant for assessing difficult moral choice situations.

Keywords Multidimensional scaling · Similarity · Conceptual spaces · Moral principle · Ethics · The ethics of technology

✉ Steven Verheyen
verheyen@essb.eur.nl

1 Introduction

The theory of conceptual spaces introduced by Gärdenfors (2000, 2014) has been applied to topics in psychology, linguistics, computer science, and philosophy (e.g., Chella et al. 2001; Cubek et al. 2015; Decock et al. 2014; Douven 2016, 2019; Douven et al. 2013; Gärdenfors and Zenker 2013; Valentine et al. 2016; Verheyen and Égré 2018; Zenker and Gärdenfors, 2015). As explained by Douven and Gärdenfors (2019), the central idea is that concepts can be modelled as regions in similarity spaces. A stock example is our perceptual color space. Different shades of blue, such as azur blue, navy blue, light blue, and ice blue are perceived as more similar to each other than, say, ice blue and mahogany red. By ordering colors according to their perceived similarity, they can be represented in a cone-shaped conceptual space along three dimensions (hue, chroma, and lightness) such that the distance between any two shades represents their degree of similarity. For instance, light blue and ice blue shades are closer to each other than ice blue and mahogany red (Indow 1988; Shepard 1964).

By modelling concepts as regions in conceptual spaces we can, among other things, explain why children are able to learn new concepts quickly (Gärdenfors 2000). Rather than memorizing necessary and sufficient conditions for each concept, children learn to recognize a small number of prototypes, which they use for categorizing new items. A penguin is, for example, categorized as a bird rather than a fish because it is judged more similar to crows and other prototypical birds than to prototypical fish. In this model, a “bird” in the child’s conceptual space is represented by the region demarcated by all objects that are more similar to prototypical birds than to any other prototypes.

Morality is a so far underexplored domain for the theory of conceptual spaces. Can moral principles be construed as conceptual spaces? If so, how? In a recent book, *The Ethics of Technology: A Geometric Analysis of Five Moral Principles*, Peterson (2017) argues that five moral principles commonly applied for assessing the pros and cons of new and existing technologies can indeed be construed as forming a conceptual space. His work can be understood as an attempt to apply some of Gärdenfors’ key ideas to ethics.¹ The basic building blocks are moral choice situations, referred to as *cases*. Peterson asked participants to make pairwise comparisons of moral similarities across a set of cases and to select which moral principle (from an open-ended list of principles) should be applied for resolving each case. On a group level, participants tended to apply the same moral principles to cases rated as morally similar. Moreover, for some cases considered being prototypical of a principle, as many as 90 % selected the same principle. These findings seem to indicate that moral principles can be construed as regions in a shared moral space. However, Shrader-Frechette (2017) and Lokhorst (2018) criticize this proposal. They claim that morality most likely displays too many individual differences to yield a single multidimensional space of moral principles. According to Shrader-Frechette (2017),

each agent likely employs her own implicit dimension(s) to answer Peterson’s moral-similarity request. Thus for the same two cases, one agent might estimate “moral similarity” with respect to catastrophic consequences, while another

¹ Gärdenfors’ (2000, 2014) work is not concerned with ethics. He does not discuss whether moral principles could qualify as geometrically construed concepts and he leaves it open how the theory should be empirically tested. Multidimensional scaling, which is the method we will use, is one of several options mentioned by Gärdenfors.

might estimate similarity with respect to fairness. If so, Peterson has a common, moral-similarity label, but no common concept. Because different agent-responses likely presuppose different moral-similarity concepts, their responses don't make logical contact. If so, there's little justification for Peterson's quantifying and aggregating many agents' moral-similarity estimates.

Peterson (2018) addresses some of the conceptual concerns raised by Shrader-Frechette and Lokhorst, but he presents no new data. The present study is designed to fill this gap in the literature. We present results from a large study ($n = 475$) indicating that moral principles can indeed be construed as regions in a shared conceptual space. We asked each participant to make 45 pair-wise comparisons of moral similarities across a set of ten cases ($10 \times 9 / 2$ comparisons). We then averaged similarity data across participants and conducted multidimensional scaling (MDS), which enabled us to construe a “moral” conceptual space. Our findings indicate that measures of group reliability are high and that there is true structure in averaged similarity data. Contrary to what Shrader-Frechette suggests, it is thus possible to use the theory of conceptual spaces for identifying shared moral principles across a group of participants, although there are also some noteworthy individual variations.

The structure of this paper is as follows. In Sections 2 and 3, we present the background and design of the study. In Section 4, we present our findings, including a discussion of the limitations of using similarity data for characterizing moral principles as regions in a shared conceptual space. Finally, in Section 5, we state our conclusions.

2 Background and Design of the Study

An appropriate point of departure for the construction of a shared conceptual space for moral principles is Aristotle's famous remark in the *Nicomachean Ethics* that moral agents should “*treat like cases alike*”.² Aristotle's principle is widely accepted by contemporary ethicists. For instance, Beauchamp and Childress (1979, 2001) agree that ethical issues encountered by medical doctors and other healthcare professionals should sometimes be analyzed by comparing how similar or dissimilar they are to cases we are already familiar with. Such comparisons help us to identify what principle(s) one ought to apply to each case. Beauchamp and Childress mention four principles they believe are applicable *prima facie* to the biomedical domain: the principle of informed consent, the principle of nonmaleficence, the principle of beneficence, and the principle of justice.

Peterson (2017) works in the same tradition as Beauchamp and Childress, but proposes a different set of principles for evaluating new and existing technologies: the cost-benefit principle (CBA), the precautionary principle (PP), the sustainability principle (ST), the autonomy principle (AUT), and the fairness principle (FP).³ Another

² NE 1131a10-b15.

³ The five principles are defined and explained in Peterson (2017). See Appendix A1 for an overview. The set is open-ended and “*could be revised if ... we were to encounter new cases that could not be plausibly analyzed by the five principles*” (Peterson 2017: 16). See Chapter 1 for a discussion of how the five principles were generated (as moral explanations of intuitions about prototype cases) and under what conditions it might be reasonable to revise the principles. This may happen if we, for instance, encounter cases that are very dissimilar from the prototypes of the five principles.

important difference is that Peterson explicitly argues that his principles can be construed as regions in a conceptual space. The key premises of this theory can be summarized as follows: If two cases are fully similar in all morally relevant aspects, then, if a principle is applicable to one case, it is also applicable to the other. Furthermore, if a case is more similar to a prototype for principle p than to the most similar prototype for any other principle, then the case should be analyzed by applying p rather than any other principle.

By identifying cases that serve as prototypes for each moral principle, the boundaries between cases covered by different principles can be represented in a Voronoi tessellation. A Voronoi tessellation divides a conceptual space into a number of regions such that each region consists of all points that are closer to a prototype for that region than to any other prototype. Within each region, the moral analysis is governed by the principle corresponding to the prototype in question. See Fig. 1.

The aim of the present study is to shed light on the empirical adequacy of the hypothesis that moral principles can be construed as regions in a conceptual space. We will focus in particular on whether morality yields a *shared* conceptual space on the *group* level, since Shrader-Frechette (2017) and Lokhorst (2018) question precisely that idea, as mentioned in the introduction.

We invited students taking a course in engineering ethics to complete an online survey. Participants were presented with ten cases (described in about 100–200 words each) featuring ethical issues related to technology and engineering. In one part of the survey, participants were asked to answer the following question: “Which moral principle should in your opinion be applied to this case?” This was followed by six options: the cost-benefit principle (CBA), the precautionary principle (PP), the sustainability principle (ST), the autonomy principle (AUT), the fairness principle (FP), and “none of the principles listed here” (see Appendix A1 for precise formulations). The ten cases were selected with the intention of identifying two paradigm cases for each principle. In the other part of the survey, participants were invited to make pairwise

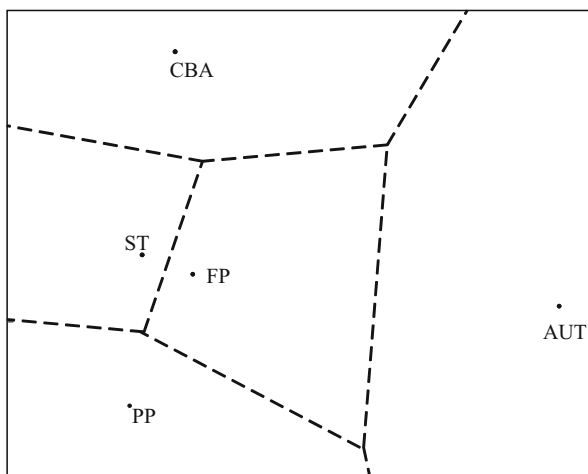


Fig. 1 An example with five moral principles represented as a two-dimensional conceptual space. A particular case is analyzed by applying the principle applicable to the most similar (nearest) prototype. See Chapter 2 in Peterson (2017) for details

similarity comparisons of all cases. For ten cases, this generated 45 pairwise comparisons, each of which was preceded by the following question: “*How similar are the following cases from a moral point of view?*”. We also varied the order between the two types of questions in the survey (“*Which moral principle ...*” and “*How similar are ...*”). In the first version (Survey A), participants were asked to make the similarity comparisons at the end of the survey; in the second version (Survey B), they were asked to make the similarity comparisons at the beginning.

Similarity data collected in both surveys were analyzed with multidimensional scaling techniques (MDS; Borg and Groenen 2005; Kruskal and Wish 1978). The term MDS refers to a family of statistical models that represent measurements of (dis)similarity among pairs of stimuli as distances between points in a low-dimensional multidimensional space. This makes it possible to uncover nonobvious structures among stimuli. Without offering instructions to participants about the characteristics on which the similarity judgments are to be made, and without having participants verbalize their considerations, the basis of their judgments can be revealed by relating geometric properties of the representation (e.g., dimensions, partitions, clusters, ...) to substantive information about the represented stimuli. In interval MDS, all pairs of stimuli i and j are positioned in space such that their distance d_{ij} corresponds to a linear transformation of their perceived similarity $f(s_{ij})$ (with smaller distances denoting greater similarity and vice versa). The distances between points i and j are normally measured by using the familiar Euclidean metric, but alternative distance functions are of course possible and sometimes more appropriate.

The extent to which the distances represented in MDS successfully capture the transformed input similarities is reflected in the squared error of each representation: $[f(s_{ij}) - d_{ij}]^2$. These discrepancies can be depicted in a Shepard diagram to determine which pairs are particularly poorly represented. A Shepard diagram contains a scatter plot of the input similarities versus the corresponding distances in the MDS space, as well as a regression line representing the optimally transformed similarities. A point's squared vertical distance from the regression line indicates the corresponding pair's residual error. The discrepancies can also be summed across all pairs in which a particular stimulus features, to establish how badly an individual point is fitted, or summed across all pairs to obtain an indication of how well the input similarities as a whole are represented. The former measure is commonly referred to as stress per point, while the latter is called stress.⁴ If these badness-of-fit indications are sufficiently low, MDS yields a visual representation of the empirical relations that exist between the stimuli, which tend to be easier to interpret than the numerical indices of these relationships.

Information about the participants providing the similarity judgments can be invoked for analyzing variations among individual responses. Individual differences scaling (INDSCAL; Carroll and Chang 1970; Takane et al. 1977) structurally incorporates individual differences by estimating individual weights for each of the dimensions of a so-called group space. By multiplying an individual's weights with the coordinates of the stimuli in the group space, one arrives at that individual's individual stimulus representation. The weights thus achieve a stretching or compression of the group space, reflecting the importance each individual attaches to different dimensions

⁴ In practice, a measure called stress-1 tends to be used, which corresponds to a normalization of the stress value by dividing it by the sum of the squared distances d_{ij} and taking the square root of the result.

of that space. The better we can approximate an individual's similarity data through a weighting of the dimensions of the group space, the lower the stress-per-person (the stress measure for that particular person) will be. If the group space shows no correspondence at all to an individual's similarity data, that will be visible in the estimated dimension weights, which will tend to be close to zero (indicating that the organization of the stimuli along the dimensions of the group space bears no resemblance to the similarity structure provided by the individual).

In the following sections, we use interval MDS as well as INDSCAL for analyzing the moral similarity judgments reported by participants. The aim is to determine whether these similarity judgments provide a sufficiently reliable basis for constructing a common space of moral principles.

3 Methods

3.1 Participants

Four hundred and seventy-five students taking a mandatory course in engineering ethics at the College Station campus of Texas A&M University's College of Engineering completed one of two versions of an online survey in exchange for partial course credit. From the 219 students who fully completed Survey A, 46 were removed (21%) because they failed at least one control question, indicated not to understand the instructions, and/or indicated that their effort was insufficient for including their data in a scientific report. A total of 173 responses to Survey A were thus retained for further analysis. From the 256 students who fully completed Survey B, 37 were removed (14%) for a total of 219 participants in Survey B. The median time spent by these participants on the survey was 25 min, 59 s.

At the request of the Institutional Review Board at Texas A&M University, no demographic information was collected. However, the demographics of the student sample that was invited to participate is publicly available at <https://accountability.tamu.edu/All-Metrics/Mixed-Metrics/Student-Demographics>. It was mainly comprised of men (78.2% versus 21.8% female). The majority of the students in the class were aged 18–21 (53.00%) or 22–25 (35.13%). The most represented ethnicities were White (46.6%), Hispanic (21.02%), International (14.58%), and Asian (11.80%).

3.2 Materials

Participants were presented with ten cases (vignettes, described in about 100–200 words each) featuring ethical issues related to technology and engineering. Each case was chosen to be representative of one of five moral principles: the cost-benefit principle (CBA), the precautionary principle (PP), the sustainability principle (ST), the autonomy principle (AUT), and the fairness principle (FP). Care was taken that the two cases deemed prototypical of a particular principle did not share apparent surface or content similarities. For instance, the two cases for the AUT principle were set in China and the US. One dealt with internet censorship and the other with fracking.

Table 1 provides an overview of the 10 cases and the five principles designed to be applicable to them. See Appendices A1 and A2 for precise definitions of each moral principle and summaries of the cases. Some cases were identical to those used by

Table 1 Overview of the 10 moral cases included in the study

Case	Presumed principle	Identifier
Prioritizing Design Improvements in Cars	Cost-Benefit	CBA1
The Cost of CO2 Emissions	Cost-Benefit	CBA2
The Challenger Disaster	Precautionary	PP1
Is Trichloroethylene a Human Carcinogen?	Precautionary	PP2
The Ten Mile Creek Basin	Sustainability	ST1
Biodiesel in the Transport Sector	Sustainability	ST2
Internet Censorship and Surveillance in China	Autonomy	AUT1
Fracking and Self-governance in Denton	Autonomy	AUT2
Broadband Access in North Dakota	Fairness	FP1
The SpaceX Launchpad in Boca Chica	Fairness	FP2

Peterson (2017), but since our aim was to identify two prototypes for each principle, a couple of new cases were developed from scratch.

3.3 Procedure

Participants completed an online survey consisting of an applicability and a similarity judgment task. There were two versions of the survey: In Survey A participants completed the applicability task before the similarity task; in Survey B participants completed these tasks in reverse order.

In the applicability task, participants were asked to answer the following question for each of the 10 moral cases: “Which moral principle should in your opinion be applied to this case?” This was followed by six answer options: the five principles (including their definition) listed in Appendix A1 and “none of the principles listed here”. We randomized the order in which the cases were presented, as well as the answer options for each case.

In the similarity judgment task, participants were invited to make pairwise similarity comparisons of all cases. For ten cases, this generated 45 pairwise comparisons, each of which was preceded by the following question: “How similar are the following cases from a moral point of view?” We offered no instructions concerning the characteristics on which these similarity judgments were to be made. We explicitly indicated that participants were NOT to make their judgments based on accidental factual, physical, or historical similarities between the cases. Participants provided their responses on a 7-point Likert scale ranging from “very dissimilar” to “very similar”. We randomized the order in which the pairs of cases were presented. Prior to the start of the applicability and similarity tasks, participants were asked to indicate whether they understood the instructions.

A key difference between the present study and that reported in Peterson (2017: Chapter 3) is that every participant in the present study was instructed to compare *all* possible combinations of all ten cases, instead of just making a small subset of such comparisons. This generated a relatively high workload for participants, so to ensure that they were paying attention we included three control questions of the type “This is

a control question to check whether you are paying attention. Please proceed by clicking 1 (Very dissimilar) on the scale below.” At the end of the survey, participants answered an additional question that read: “Have you answered the questions to the best of your ability? Do you feel that your effort is sufficient for including your data in a scientific research report? Please be honest. You will receive the extra credit regardless of how you answer this question.” Participants answered by either clicking “Yes, I have answered the questions to the best of my ability. Please include my answers in your study.” or “No, I think my answers should be omitted, but I will receive the extra credit anyway.”

4 Results

4.1 Reliability

We determined the reliability of the similarity judgments by applying the Spearman-Brown formula to the split-half correlations (Spearman 1904). The reliability of similarity data in both Survey A and Survey B was established at .99. The reliability remained at a high value of .97 for both Survey A and B when the data were split in half (first half of participants, second half of participants, even participants, uneven participants) so this does not appear to be an artefact of simply having a large number of participants provide the similarity judgments. If we restrict the sample size to 10% of the original samples (17 participants for Survey A and 22 participants for Survey B) and calculate the reliability for 1000 such samples, we still get average reliabilities of .88 and .86, respectively.

Similarity data in Survey A and B were averaged and transformed to dissimilarities by subtracting the average similarity for each pair from 8 (the maximum similarity scale value plus one). The resulting dissimilarities were subjected to interval multidimensional scaling using the smacof package (De Leeuw and Mair 2009) in R version 3.6.1 (R Core Team 2017). Following Peterson (2017), we obtained solutions in two and three dimensions. The resulting stress-1 values were .155 and .088 for Survey A, and .167 and .088 for Survey B. These empirical badness-of-fit values are lower than the stress-1 values obtained for random input dissimilarities. The average stress-1 value across 10,000 simulated data sets comprising 10 by 10 dissimilarities randomly sampled from a uniform distribution between 0 and 1 equals .235 for two-dimensional MDS configurations and .140 for three-dimensional MDS configurations with standard deviations of .020 and .016, respectively. The empirical stress-1 values fall just below the critical values of .174 (in 2D) and .092 (in 3D) obtained by taking the mean stress-1 value for the random data minus 3 times the standard deviation (Spence and Ogilvie 1973).

The MDS configurations for Survey A as well as Survey B also pass a permutation test in which the empirical stress-1 values are compared to a distribution of stress-1 values obtained by subjecting permutations of the empirical input (dis) similarities to MDS (Mair et al. 2016). An advantage of this procedure (over the comparison with randomly generated dissimilarities discussed in the previous paragraph) is that it respects the nature of the data. It yields a test of the assumption that the input (dis) similarities are interchangeable. Assuming an $\alpha = .05$, this null hypothesis is rejected

for the Survey A data in two ($p = .0003$) and three dimensions ($p = .0007$) as well as for the Survey B data ($p = .0009$ and $p = .0008$, respectively).⁵

4.2 Organization of the Moral Spaces

The upper panels of Figs. 2 and 3 show two- and three-dimensional MDS representations of the averaged similarity data in Survey A (left) and Survey B (right). Note that the 10 moral cases are similarly organized in both MDS configurations, in two as well as in three dimensions. This indicates that there is common structure underlying these representations. If they had been based on random or widely diverging similarity judgments, the 10 cases would almost certainly have been positioned differently in different spaces. Therefore, the observation that the MDS configurations for Survey A and Survey B yield similar structures signals that the same considerations informed the underlying similarity judgments. It also indicates that there is no apparent effect of having participants apply the moral principles to the cases before (Survey A) or after (Survey B) making similarity judgments. Regardless of whether participants were primed or not with the five moral principles, they appear to perceive the 10 cases in a similar way.

The MDS configurations also show that the two cases that were presumed to be prototypical for each of the included moral principles (see Table 1) cluster together in space. This indicates that the principles that were used for the construction and selection of the cases informed the participants' similarity judgments. There is no reason for these cases to end up so close together in space unless participants picked up on these commonalities in their assessments of moral similarities.⁶ Because cases that are prototypical for a principle cluster together in similarity space, they can be used to carve out regions in space that denote the five moral principles. One can thus identify a Voronoi tessellation of the similarity space, in which each Voronoi cell is comprised of those points in space that are closest to the two prototypical instances of each principle. That is, one can conceive of the similarity space as a "moral" conceptual space representing the cost-benefit principle (CBA), the precautionary principle (PP), the sustainability principle (ST), the autonomy principle (AUT), and the fairness principle (FP), each represented by two prototypes. Note that this is only formally achievable in the three-dimensional configurations (Fig. 3). In the two-dimensional configurations (Fig. 2), the AUT and FP cases cannot be clearly discerned; doing so requires the addition of a third dimension. Mirroring the close similarity between the AUT and FP cases, the fairness principle was often applied to the autonomy cases and vice versa (Table 2 shows the number of times each principle was applied to the 10 cases). We will return to the implications of this finding in Section 4.5.

4.3 Individual Differences

Although the average similarity data for surveys A and B are reliable, there are some noteworthy individual differences among participants. The average correlation between

⁵ Similar results were obtained with non-metric MDS. In non-metric MDS, the transformation of the perceived similarities $f(s_{ij})$ is monotonic instead of linear as in interval MDS.

⁶ As noted earlier, we took care to limit surface similarities between cases belonging to the same moral principle, and we also instructed participants to disregard accidental factual, physical, or historical similarities between cases.

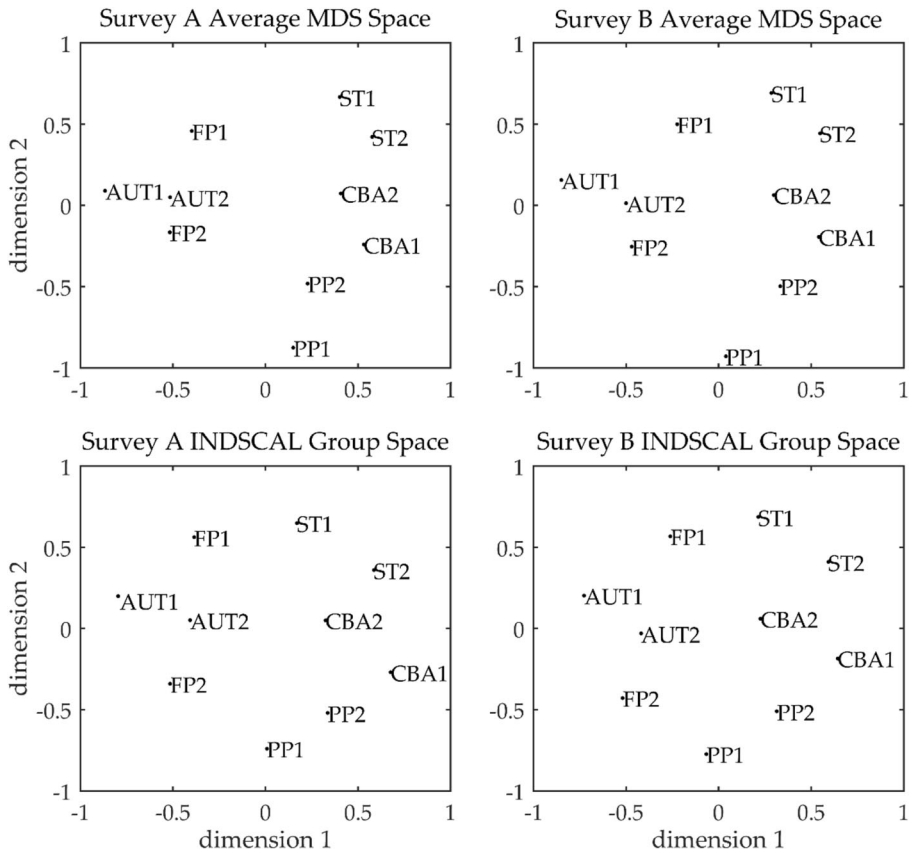


Fig. 2 Two-dimensional configurations of the Survey A (left) and Survey B (right) similarity data. The upper panels represent the configurations obtained for the average similarity data using regular interval MDS. The lower panels represent the group configurations obtained with individual differences scaling (INDSCAL). The configurations were brought into the same orientation using Procrustes analysis (Gower and Dijksterhuis 2004)

participants' similarity judgments is only .29 for Survey A and .23 for Survey B. This considerable inter-individual variability is mirrored in the variability found in the applicability judgments (i.e., the responses to the question about which principle should be applied to each case). Fleiss's (1971) kappa for the applicability data in Survey A is .44 and for Survey B it is .47, which is about halfway between perfect agreement and agreement due to chance. Only 9% of participants (8% in Survey A and 10% in Survey B) classified all 10 cases as we had intended. The prototypical cases are, however, identified as prime examples of the five moral principles, as is shown in Table 2. The principle we had intended to be chosen was the predominantly chosen principle for each of the 10 cases, both in Survey A and B. However, the applicability percentages in Table 2 indicate that the cases designed to be prototypical for the cost-benefit and the fairness principles were considered less prototypical than the cases for the other principles. The precautionary principle was often found to apply to CBA1, while the sustainability principle was often found to apply to CBA2. Many participants found the cost-benefit principle to also apply to FP1 and the autonomy principle to FP2.

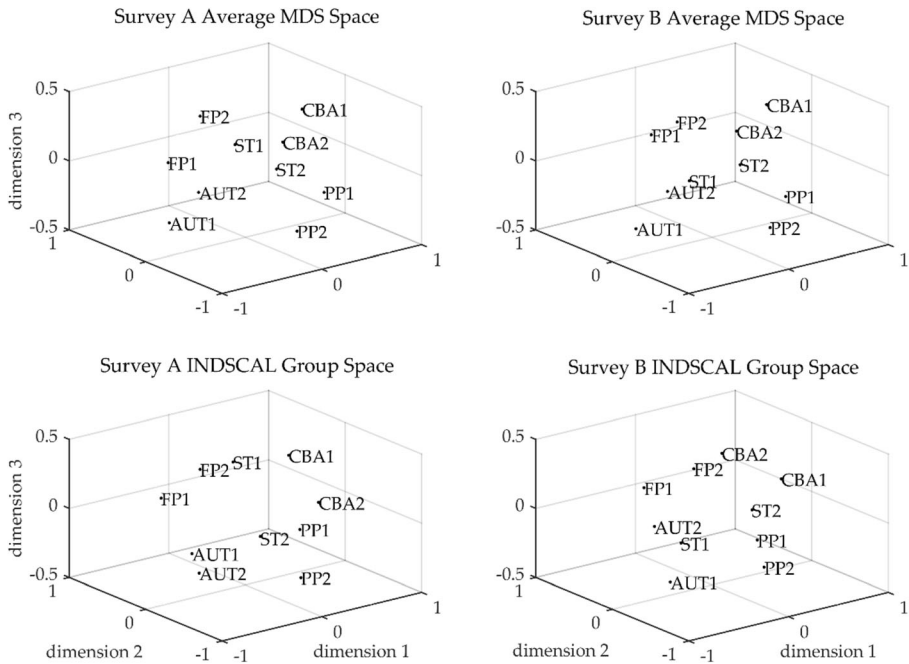


Fig. 3 Three-dimensional configurations of the Survey A (left) and Survey B (right) similarity data. The upper panels represent the configurations obtained for the average similarity data using regular interval MDS. The lower panels represent the group configurations obtained with individual differences scaling (INDSCAL). The configurations were brought into the same orientation using Procrustes analysis (Gower and Dijksterhuis 2004)

Participants in Survey A as well as in Survey B almost always judged at least one of the five moral principles to be applicable to the ten moral cases. The option “*none of the*

Table 2 Principle applicability percentages for each of the 10 moral cases in Surveys A and B

	Cost-Benefit		Precautionary		Sustainability		Autonomy		Fairness		None	
	A	B	A	B	A	B	A	B	A	B	A	B
CBA1	.58	.68	.27	.22	.02	.01	.05	.01	.04	.02	.03	.05
CBA2	.46	.55	.17	.12	.33	.27	.02	.01	.01	.02	.00	.02
PP1	.03	.02	.86	.82	.00	.00	.03	.03	.02	.01	.06	.11
PP2	.02	.04	.86	.79	.06	.07	.02	.05	.03	.01	.02	.03
ST1	.13	.10	.03	.03	.75	.76	.02	.02	.03	.04	.04	.05
ST2	.15	.14	.06	.06	.73	.76	.01	.01	.01	.00	.03	.03
AUT1	.01	.00	.03	.04	.01	.01	.83	.81	.10	.11	.02	.03
AUT2	.05	.01	.04	.03	.12	.08	.69	.78	.08	.10	.03	.00
FP1	.19	.14	.01	.00	.03	.02	.08	.12	.65	.66	.05	.06
FP2	.05	.08	.06	.05	.05	.02	.17	.23	.62	.59	.05	.03
Total	.17	.18	.24	.22	.21	.20	.19	.21	.16	.16	.03	.04

principles listed here” was chosen in no more than 3% of the 10×173 trials in Survey A and in 4% of the 10×219 trials in Survey B.

We performed individual differences scaling (INDSCAL; Carroll and Chang 1970; Takane et al. 1977) on the similarity data for Survey A and Survey B, which yielded group configurations (lower panels of Figs. 2 and 3) that were very similar to the configurations of the average data (upper panels of Figs. 2 and 3).⁷ The INDSCAL analyses yielded no individuals with weights close to zero, which would have been an indication that those individuals’ data did not line up very well with the group space. The minimum individual weights for data in Survey A when analyzed in two dimensions were .82 and .99, and in three dimensions .76, .93, and 1.25. The minimum individual weights for data in Survey B when analyzed in two dimensions were .88 and .93, and in three dimensions .86, .89, and 1.15.

The INDSCAL analyses also yielded a quantitative indication of how well a participant matches the group configuration: the stress-per-person is lower the better that person’s data can be obtained through a weighting of the group configuration’s dimensions. We established a positive correlation between stress-per-person and the number of misclassifications of a person (identified as the number of times out of 10 that the person selected a different principle for a case than the one we intended). For Survey A, we established Pearson’s linear correlation coefficient at .31 ($p < .0001$) in both two and three dimensions. For Survey B, these values measured .26 ($p = .0001$) and .35 ($p < .0001$), respectively. These correlations suggest that the more one believes other principles apply to the moral cases, the less one’s similarity configuration fits the group organization in terms of the 5×2 prototypes.

4.4 Boundary Conditions

The observation that we can identify a common structure in averaged similarity judgments only holds when two prototypes per principle are included. From the similarity data in both Survey A and Survey B one can construct 32 different data sets with one prototype per principle if one considers all possible combinations of prototypical cases. When these data sets were subjected to interval MDS, the large majority failed the permutation test at $\alpha = .05$, both for Survey A and for Survey B and in two and three dimensions.⁸ When only one prototypical case is included per principle, the null hypothesis that the input (dis) similarities are interchangeable thus cannot be rejected.⁹ This suggests that the structure we established in the spaces with two prototypical cases per moral principle is mostly local. It appears to derive primarily from the high similarity of the two prototypes for each principle. This observation is corroborated by the Shepard diagrams (not shown) that indicate that the smaller dissimilarities are better captured by the MDS distances than the larger dissimilarities.

⁷ One participant provided identical similarity judgments for all pairs in Survey B. This participant was excluded from the INDSCAL analysis to overcome technical issues.

⁸ The stress test is not well-defined for data sets with a small number of items (here: 5) and therefore not considered here.

⁹ A similar conclusion was reached based on the analyses of similarity judgments obtained for cases CBA1, PPI, ST1, AUT1, and FPI in a sample of 204 students drawn from the same course in engineering ethics (none of whom had participated in Survey A or B).

The structure that is present in the MDS configurations of the entire set of cases appears to largely reflect the applicability of the five moral principles. For each participant, we constructed an alternative similarity matrix based on the principles they applied to each of the cases. Pairs of cases that were awarded the same principle received a similarity score of 1; pairs of cases that were awarded different principles a similarity score of 0. (This procedure corresponds to the free sorting procedure for obtaining similarity data used in many MDS applications, where participants sort stimuli into piles with the understanding that stimuli in the same pile have something in common, while stimuli in different piles do not; Borg and Groenen 2005; Miller 1969). The individual similarity matrices were averaged across participants, subsequently transformed to a dissimilarity matrix by subtracting the average similarity for each pair from 1, which was then subjected to interval MDS. The resulting three-dimensional configurations for the Survey A applicabilities (left) and the Survey B applicabilities (right) are shown in Fig. 4. They closely resemble the configurations in Fig. 3, both in terms of the clustering of the two prototypical cases per principle, but also in terms of the overall structure of the configuration. This suggests that the larger distances in the original, similarity-based configurations tend to capture some of the systematic differences in opinion as to whether which principles apply to the cases, in that cases that are seldomly awarded the same principle are also judged to be less similar.

Because of the pronounced inter-individual differences, a certain number of participants is required to obtain reliable MDS configurations. While the two- and three-dimensional configurations obtained on half of the similarity data (first half of participants, second half of participants, even participants, uneven participants) all pass both the stress and permutation tests, a considerable number of configurations fail these tests when they are based on similarity data from a smaller sample. We randomly drew 100 samples of varying sizes of similarity data from Survey A and Survey B and subjected the averaged data from each sample to interval MDS. With samples sizes of 20, 41% of Survey A samples and 30% of Survey B samples failed at least one of the tests in two dimensions. The corresponding percentages in three dimensions were 50% and 37%. Although the reliability of such samples is quite high (see section 4.1 for the average reliabilities for samples corresponding to 10% of the sample $\approx N = 20$), these results signal the need to conduct MDS-specific tests to assess whether the resulting MDS

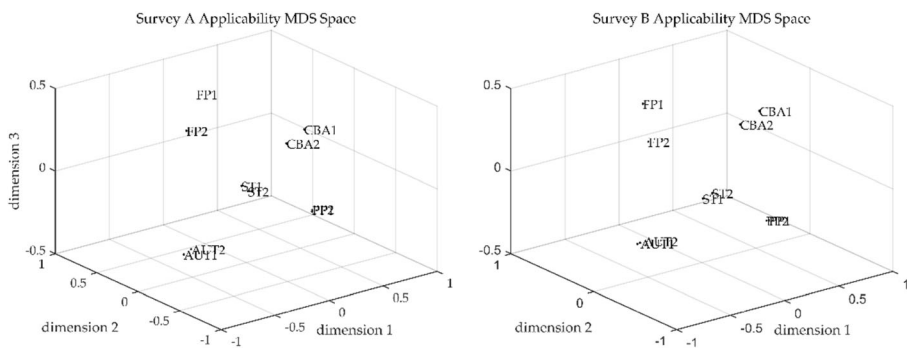


Fig. 4 Three-dimensional configurations of the Survey A (left) and Survey B (right) applicability data. The number of times pairs of cases were awarded the same moral principle was subjected to regular interval MDS. The configurations were brought into the same orientation as the ones in Fig. 3 using Procrustes analysis (Gower and Dijksterhuis 2004)

configurations should be interpreted. With sample sizes of 40, the percentage of samples that failed at least one of the tests halved, to 19% and 14% in two dimensions, and 28% and 9% in three dimensions for Survey A and Survey B, respectively. With a sample size of 80 (nearing half of our original sample sizes), almost all samples passed both tests. The corresponding percentages of samples failing at least one test were 6%, 3%, 14%, and 1%.

4.5 Key Findings

The most important finding in light of the criticism voiced by Shrader-Frechette (2017) and Lokhorst (2018) is that there is common structure to be found in averaged similarity judgments of moral choice situations. The high reliability measures indicate that the averaged similarity judgments are stable across groups. This is a requirement for the MDS configurations to be representative of a structure shared among participants (Ashby et al. 1994; Lee and Pope 2003) and for the configurations to be replicable (Sturidsson et al. 2006; Voorspoels et al. 2014; White et al. 2014). The stress tests and permutation tests conducted on the MDS configurations of the average similarity data indicate that there is structure underlying these configurations. The data generation process triggering the similarity judgments is thus neither random or completely idiosyncratic (Spence and Ogilvie 1973; see also Klahr 1969, Stenson and Knoll 1969, and Sturrock and Rocha 2000), nor are the similarity judgments of different moral cases interchangeable (Mair et al. 2016).

The three-dimensional MDS configurations indicate that the two cases deemed prototypical of each principle cluster together. This finding supports Peterson's (2017) claim that by applying MDS to similarity judgments of moral choice situations one can construct conceptual spaces in which moral principles are discernable as distinct regions. Moreover, by assessing the similarity of new moral cases to the ones that are prototypical for the five principles, it is possible to determine which principle to apply when assessing new cases (see Chapter 3 in Peterson 2017, for an illustration). However, in the two-dimensional MDS representations this structure did not come out as expected. As noted, the cases representing the autonomy and fairness principles could not be discerned in two-dimensional representations. One might perhaps argue that this was due to those representations lacking one of the three dimensions that differentiate the principles. However, one could also take this observation to be a reason for preferring a more parsimonious space with, for instance, four instead of five principles.¹⁰ Conversely, one can also imagine enriching the space by adding cases deemed to be prototypical of other moral principles and study whether those cases are sufficiently different from the ones already present in the space.¹¹

¹⁰ It is worth keeping in mind that it may be no easy task to formulate a single principle that articulates intuitions about fairness *and* autonomy, but not any of the other values alluded to in the other principles.

¹¹ What constitutes a sufficient difference or a sufficient similarity is another question, which we do not attempt to answer here, but can be addressed through clustering. See Verbeemen, Vanpaemel, Pattyn, Storms, and Verguts (2007) for an example. An alternative approach might be to recognize explicitly that cases have both overlapping and distinctive features and represent the similarities through additive cluster models (Shepard and Arabie 1979) or extended trees (Corter and Tversky 1986) instead of geometric spaces.

The observation that the cost-benefit principle (CBA), the precautionary principle (PP), the sustainability principle (ST), the autonomy principle (AUT), and the fairness principle (FP) can be discerned in a constructed moral space for participants without previous exposure to the principles (in Survey B the applicability task was preceded by the similarity judgment task) speaks to the relevance of these principles. Similarity judgments of moral cases and their representation in multidimensional spaces can thus help us identify the moral principles that are relevant for assessing technological innovations.

We note that these findings hold when the data is split in half. However, when data of fewer participants is used, the observed structure begins to break down. This is due to individual differences among participants. To obtain a stable, reliable structure one needs to obtain similarity judgments from a sufficiently large number of individuals, and the similarity judgments must not be heavily influenced by individuals whose opinions deviate strongly from the majority. We observed pronounced individual differences with respect to both the applicability and the similarity tasks. The results of the individual differences scaling suggest a relationship between the two: the more a participant feels that individual cases should be judged along different principles (rather than the ones intended), the less the participant agrees on the general configuration depicting the intended structure. This highlights a clear avenue for future research, namely to further investigate the origin and nature of these individual differences.

Another noteworthy finding is that several prototypes per principle have to be included to delineate all principles clearly. This is not problematic from a theoretical point of view. Several theorists working on conceptual spaces have argued for the importance of using multiple prototypes per concept, or regions of prototypical instances, per concept (e.g., Douven et al. 2013; Gärdenfors and Williams 2001; Regier et al. 2005; Storms et al. 2000). The decision to include several prototypes per principle may also have some additional advantages in moral contexts: If one conceives of the boundaries of a principle as the points that are equidistant between unique prototypes, then all boundaries between principles will be sharp. This corresponds to a moral theory in which a single moral principle governs (the perception of) a moral case. Beauchamp and Childress (1979, 2001), Peterson (2017), and others stress that it is often appropriate to apply more than one principle to a case. By using several prototypes per principle, we can model this plausible idea: Instead of having a unique delineation of the similarity space based on individual prototypes, we can produce multiple delineations based on the combinations of different prototypes per principle. Some cases will fall within the region covered by one principle on one delineation, but in the region of another principle on another delineation. The proportion of times a case falls under a particular principle (i.e., is found to be more similar to a chosen prototype of one principle than to a choice of prototypes representing other principles) can be used as a rough measure of the extent to which the principle applies (Decock and Douven 2014; Douven 2016; Douven et al. 2017; Verheyen and Égré 2018. See also Peterson 2017: Chapter 2). This allows for borderline cases to which more than one principle applies. Consider, for instance, the description of the Challenger Disaster in Appendix A2. If we were to include quantitative information about the costs of postponing the launch of the shuttle, it seems likely that this case would be placed in a gray area in which both the cost-benefit and precautionary principles apply.

5 Discussion

Our findings indicate that five moral principles frequently applied for analyzing ethical issues related to technology and engineering can be represented as regions in a shared moral space. Although we found noteworthy individual differences among participants, averaged similarity judgments of moral choice situations display a common and stable structure, contrary to the intuitions voiced by Shrader-Frechette (2017) and Lokhorst (2018). It seems likely that parallel representations in other domains of (applied) ethics are also possible. We would, for instance, not be surprised if the four principles proposed by Beauchamp and Childress (1979, 2001) for the biomedical domain (the principle of informed consent, the principle of nonmaleficence, the principle of beneficence, and the principle of justice) could be similarly represented in a shared moral space. If so, it would be interesting to investigate whether Beauchamp and Childress' moral space could be integrated with that for technology and engineering, as some borderline cases seem to belong to both (e.g., the development of new drugs). It is beyond the scope of this paper to investigate this here. However, future research may show if different domains of applied ethics can be subsumed in a higher dimensional space and whether those dimensions are integral or separable.

That said, we are of course aware that not every moral theorist will welcome our approach, for several reasons. To begin with, it might be objected that it is a mistake to develop *several* moral principles. All we need is a *single* principle that covers all cases. For instance, John Stuart Mill (1865) famously claims that an act is right just in case it maximizes overall utility, and Kant (1785) argues that an act is right only as long as it does not violate his categorical imperative.¹² We agree with Mill and Kant that unary accounts of morality are elegant and attractive from a theoretical point of view, but we insist that no single principle can explain the *descriptive* findings reported in this paper. If a single principle governs people's similarity judgments, then all cases in which, say, the categorical imperative was perceived as satisfied would have been judged fully similar to each other. Moreover, cases in which the categorical imperative was believed violated, would have been rated as maximally dissimilar to cases in which that is not the case. However, as noted in Section 4, we did not observe this type of binary or highly polarized similarity judgments. Moral outlooks that include several principles offer a better fit with our findings.

Another worry moral theorists may voice concerns the somewhat inflexible nature of the principles generated from similarity judgments. In Ross's (1930) well-known theory of prima facie duties, each of his seven principles is valid only in so far as it is not overruled by another principle. Ross claims that in order to determine what one ought to do all things considered, all prima facie principles have to be balanced against each other. This dynamic process will eventually enable the agent to identify his duty proper. However, the five principles we generate from similarity data do not seem to allow for this type of balancing of conflicting duties or values, meaning that they are more inflexible than Ross's

¹² Kant proposed several formulations of the categorical imperative. The perhaps most well-known one is the following: "Act only according to that maxim whereby you can at the same time will that it should become a universal law."

principles. Our response is that some flexibility can be achieved in our model by letting the boundaries of each principle be defined by more than one prototypical case, as noted in Section 4. If each principle is represented by several prototypes, then the regions covered by the principles will overlap. We admit that the balancing process itself is not captured by our account; the similarity judgments describe the situation *after* the balancing process. Therefore, although our account does not capture *all* aspects of Ross's famous theory, we believe it is compatible with some of its most important features.

At no point in this paper have we attempted to derive an "ought" from an "is". We accept what moral philosophers call Hume's Law, meaning that we do not believe it is possible to derive any normative recommendations from purely descriptive premises. Our aim is to study the moral opinions people actually hold; we are not making any claim about what opinions one ought to hold. However, we nevertheless believe these descriptive findings are relevant, in indirect ways, for addressing normative issues. First, our model makes it possible to check whether a set of moral judgments is internally coherent in the following sense: Do agents apply the same moral principle to cases they believe to be similar from a moral point of view? If some cases that are judged similar (meaning that they are located in the same region of moral space) were not treated alike, then those judgments would violate Aristotle's dictum that we should "*treat like cases alike*". Second, if we believe that peoples' similarity judgments are, on average, reasonably accurate, then we can analyze new cases not included in our study by comparing how similar they are to the prototypical cases we already know how to analyze. The premise that bridges the gap between "is" and "ought" here, is the assumption that peoples' similarity judgments are reasonably accurate.

Our final comment concerns the possibility of applying the methodology outlined in this paper to legal issues. In the common law tradition, judges routinely compare how similar or dissimilar new cases are to cases ruled on in previous court rulings. The normative assumption underpinning this is, again, Aristotle's insight that judges should "*treat like cases alike*". We note that our approach could be used for constructing legal similarity spaces that are analogous to the moral spaces constructed in this paper. By measuring legal similarities and dissimilarities across a set of legal cases, one could map the corresponding cases onto a multidimensional legal space. One could then verify whether cases located close to each other are treated alike, and perhaps identify the legal principle(s) applied to each case. If it transpires that cases that legal experts (or law students, or lay people) perceive as similar from a legal point of view are not treated alike, this could be a reason for questioning the underlying court rulings. This indicates that the approach to normative reasoning outlined in this paper can be applied to a fairly broad domain of issues.

Compliance with ethical standards

Ethics and Consent This study was conducted with the approval of the Human Research Protection Program at Texas A&M University (IRB ID IRB2019-0238M). Informed consent was obtained from all participants.

Open Practices Statement The data, materials, and R scripts for conducting the multidimensional scaling analyses are available at <https://osf.io/cgq4d/>. The data, materials, and scripts used in this article are licensed

under a Creative Commons Attribution 4.0 International License (CC-BY 4.0). To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

Author Note SV and MP conceptualized the study and drafted the manuscript. MP gathered the data. SV analyzed the data. All authors discussed the findings thoroughly, read, and approved the final version of the manuscript. SV was funded by KU Leuven Research Council grant C14/16032. We thank Gert Storms and James Hampton for helpful suggestions.

Appendix

Appendix A1: Definition of the 5 moral principles after Peterson (2017)

- The Cost-Benefit Principle (CBA): A technological intervention is morally right only if the net surplus of benefits over costs for all those affected is greater than that of every alternative.
- The Precautionary Principle (PP): A technological intervention is morally right only if reasonable precautionary measures are taken to safeguard against non-negligible but uncertain threats.
- The Sustainability Principle (ST): A technological intervention is morally right only if it does not lead to any significant long-term depletion of natural, social or economic resources.
- The Autonomy Principle (AUT): A technological intervention is morally right only if it does not reduce the independence, self-governance or freedom of the people affected by it.
- The Fairness Principle (FP): A technological intervention is morally right only if it does not lead to unfair inequalities among the people affected by it.

Appendix A2: Overview of the 10 moral cases

Cost-Benefit Principle cases

CBA1: Prioritizing Design Improvements in Cars In the mid 1990's it was proposed that various automobile design improvements should be made mandatory by the U.S. National Highway Traffic Safety Administration (NHTSA). According to Tengs et al. (1995) the cost of implementing each design improvement, expressed in dollars per life-year saved, would have varied between \$0 and \$450,000. (See the list below.) The benefit of each design improvement listed below is the prevention of one statistical death per year. Granting a limited budget, would it be morally right to prioritize the most cost-effective safety improvement (installing windshields with adhesive bonding) over other costlier options?

- Install windshields with adhesive bonding \$0.
- Automobile dummy acceleration (vs. side door strength) tests \$63,000.
- Collapsible (vs. traditional) steering columns in cars \$67,000.
- Front disk (vs. drum) brakes in cars \$240,000.
- Dual master cylinder braking system in cars \$450,000.

CBA2: The Cost of CO₂ Emissions. The Scientific American reports that under the Obama administration, multiple agencies working together came to the conclusion that emissions of CO₂ in the U.S. cost \$121 billion in damage per year. When factoring in how much CO₂ is emitted, this means that every ton of carbon dioxide emitted into the atmosphere costs the U.S. about \$21. An anonymous official working for the Environmental Protection Agency said “\$21 doesn’t really justify much” and estimates from U.S. Global Change Research Program indicate that only an assessment of between \$36 and \$88 per ton would justify the costs of doing what is necessary to significantly reduce CO₂ emissions. Thus, those wanting to curb CO₂ emissions must find ways to assess more damage if they wish to justify more regulation. Would it be morally right to spend tax money on reducing CO₂ emissions if the cost significantly exceeds the value of the damage prevented by the regulation?

Precautionary Principle cases

PP1: The Challenger Disaster. On January 28, 1986, the Challenger space shuttle exploded shortly after take-off from Cape Canaveral, killing its crew of seven astronauts. The cause of the explosion was a leaking O-ring in a fuel tank, which could not cope with the unusually low temperature at the day of the take-off. About six months before take-off, engineer Roger Boisjoly at Morton Thiokol, the company responsible for the fuel tank, had written a memo in which he warned that low temperatures could cause a leak in the O-ring: “The result would be a catastrophe of the highest order — loss of human life”. However, Boisjoly was unable to back up his claim with data. His point was that for all they knew a leak in the O-ring could cause an explosion, but there was no or little data to confirm or refute that suspicion. The night before the launch Boisjoly reiterated his warning to his superiors. Was it morally right for Boisjoly’s superiors to ignore his unproven warning?

PP2: Is Trichloroethylene a Human Carcinogen? Trichloroethylene is a clear non-flammable liquid commonly used as a solvent for a variety of organic materials. It was first introduced in the 1920s and widely used for a variety of purposes until the 1970s, when suspicions arose that trichloroethylene could be toxic. A number of scientific studies of trichloroethylene were initiated and in the 1990s researchers at the U.S. National Cancer Institute showed that trichloroethylene is carcinogenic in animals, but there was no consensus on whether it was also a human carcinogen. In 2011 the U.S. National Toxicology Program’s 12th Report on Carcinogens concluded that trichloroethylene can be “reasonably anticipated to be a human carcinogen”. Would it have been morally right to ban trichloroethylene for use as a solvent for organic materials in the 1990s?

Sustainability Principle cases

ST1: The Ten Mile Creek Basin The Ten Mile Creek project is part of the larger Everglades Restoration Project, which covers 16 counties over an 18,000 square-mile area. The project seeks to restore, protect and preserve the water resources of central

and southern Florida. By capturing freshwater from Ten Mile Creek and storing it during the rainy season, the amount of freshwater and sediment entering waterways can be controlled. Construction consisted of a 6000 acre-feet above ground reservoir; a pump station; a gated-water control structure for moderating the release of water back into the creek; a gated gravity control structure for draining the facility for maintenance purposes; and control structures between the deep water storage area and appurtenant structures. In addition to the obvious environmental benefits of the project, St. Lucie County will use part of the site as a nature preserve area to promote hiking, fishing, bird watching and other outdoor activities. Was it morally right to build the Ten Mile Creek Basin? (Source: Skanska Group; quoted verbatim except the last sentence).

ST2: Biodiesel in the Transport Sector An in-depth study by Sandia National Laboratories and General Motors Corporation has found that plant and forestry waste and dedicated energy crops could replace nearly a third of gasoline use by the year 2030. Using a newly developed tool known as the Biofuels Deployment Model, or BDM, Sandia researchers determined that 21 billion gallons of cellulosic ethanol could be produced per year by 2022 without displacing current crops. The study, which focused only on starch-based and cellulosic ethanol, found that an increase to 90 billion gallons of ethanol could be sustainably achieved by 2030 within real-world economic and environmental parameters. Given these findings, would it be morally right to ramp up the production of cellulosic ethanol and other biofuels? (Source: Machinery lubrication newsletter; quoted verbatim except the last sentence).

Autonomy Principle cases

AUT1: Internet Censorship and Surveillance in China The Great Firewall Project is an internet censorship and surveillance project in China controlled by the ruling Communist Party. By using methods such as IP blocking, DNS filtering and redirection, and URL filtering, the Chinese Ministry of Public Security is able to block and filter access to information deemed to be politically dissident or “inappropriate” for other reasons. As part of this policy, searching with all Google search engines was banned in Mainland China on March 30, 2010. The project started in 1998 and is still in operation. Is it morally right to censor the internet in ways that limit the freedom of internet users in China to access information deemed to be politically dissident?

AUT2: Fracking and Self-governance in Denton In the early 2010’s many cities across the United States passed ordinances restricting the areas in which fracking could occur within city limits. Some concerns were public health and depreciation of private property near the fracking sites. In November 2014 the City of Denton passed an extremely restrictive ordinance that effectively banned fracking. In March 2015 the Texas legislature responded by passing a bill that effectively took away the power of cities to regulate fracking. Property owners in Denton were upset because they felt that their independence, self-governance and freedom was restricted by this. Was it morally right to restrict the independence, self-governance and freedom of property owners in Denton wishing to ban fracking in the city?

Fairness Principle cases

FP1: Broadband Access in North Dakota About 14 million rural Americans live without broadband internet access. Fast internet can be crucial to educational and economic success. “For rural residents,” writes Sharon Stover, a communications professor at University of Texas-Austin, “having broadband is simply treading water or keeping up. Not having it means sinking.” People living in the city might take for granted their ability to take online classes, navigate websites like Healthcare.gov, or apply for jobs in other places. Because the sparse populations of places like North Dakota make it expensive to pay private companies to bring high-speed internet to rural regions, some officials have called for a public solution to the problem of rural Americans being disadvantaged by lack of broadband access. Given that so many rural Americans are disadvantaged by lack of broadband access, it has been suggested that society should act to ensure that all Americans can benefit from this new technology regardless of location. Would it be morally right to subsidize broadband access in North Dakota?

FP2: The SpaceX Launchpad in Boca Chica In August 2014, Elon Musk and Rick Perry announced that SpaceX would build the world’s first private rocket launchpad on a remote beach not far from the city of Brownsville. The local government offered SpaceX more than 15 million dollars in incentives and The Greater Brownsville Incentives Corporation offered an additional 5 million in incentives to lure SpaceX away from sites in Florida and Georgia. Located only two miles from the launchpad, residents of the small village of Boca Chica will be subject to evacuations during launches and their property is at risk of damage from explosions and dangerous chemicals even during successful operations. Residents of the larger nearby city of Brownsville will reap interest on their financial investment to bring SpaceX to the area, but residents of Boca Chica are skeptical that they will receive any compensation for their much greater sacrifice. Is it morally right for SpaceX to not compensate the residents of Boca Chica?

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ashby, F.G., W.T. Maddox, and M.D. Lee. 1994. On the dangers of averaging across subjects when using multidimensional scaling or the similarity-choice model. *Psychological Science* 5: 144–151.
- Beauchamp, T.L., and J.F. Childress. 1979. *Principles of biomedical ethics*. New York, NY: Oxford University Press.
- Beauchamp, T.L., and J.F. Childress. 2001. *Principles of biomedical ethics*. 5th ed. New York, NY: Oxford University Press.

- Borg, I., and P.J.F. Groenen. 2005. *Modern multidimensional scaling: Theory and applications*. New York, NY: Springer. <https://doi.org/10.1007/0-387-28981-x>.
- Carroll, J.D., and J.J. Chang. 1970. Analysis of individual differences in multidimensional scaling via an N-way generalization of Eckart-Young decomposition. *Psychometrika* 35: 283–319.
- Chella, A., M. Frixione, and S. Gaglio. 2001. Conceptual spaces for computer vision representations. *Artificial Intelligence Review* 16: 137–152.
- Corter, J.E., and A. Tversky. 1986. Extended similarity trees. *Psychometrika* 51: 429–451.
- Cubek, R., Ertel, W., & Palm, G. (2015). High-level learning from demonstration with conceptual spaces and subspace clustering. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 2592–2597). IEEE.
- Decock, L., and I. Douven. 2014. What is graded membership? *Noûs* 48: 653–682.
- Decock, L., I. Douven, C. Kelp, and S. Wenmackers. 2014. Knowledge and approximate knowledge. *Erkenntnis* 79: 1129–1150.
- De Leeuw, J., and P. Mair. 2009. Multidimensional scaling using majorization: SMACOF in R. *Journal of Statistical Software* 31: 1–30.
- Douven, I. 2016. Vagueness, graded membership, and conceptual spaces. *Cognition* 151: 80–95.
- Douven, I. 2019. Implicatures and naturalness. In *Cognitive structures: Linguistic, philosophical and psychological perspectives*, ed. S. Löbner, T. Gamerschlag, T. Kalenscher, M. Schrenk, and H. Zeevat. New York, NY: Springer in press.
- Douven, I., L. Decock, R. Dietz, and P. Egré. 2013. Vagueness: A conceptual spaces approach. *Journal of Philosophical Logic* 42: 137–160.
- Douven, I., and P. Gärdenfors. 2019. *What are natural concepts?* *Mind & Language: A design perspective* in press.
- Douven, I., S. Wenmackers, Y. Jaissati, and L. Decock. 2017. Measuring graded membership: The case of color. *Cognitive Science* 41: 686–722.
- Fleiss, J.L. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin* 76: 378–382.
- Gärdenfors, P. 2000. *Conceptual spaces: The geometry of thought*. Cambridge, MA: MIT Press.
- Gärdenfors, P. 2014. *The geometry of meaning: Semantics based on conceptual spaces*. Cambridge, MA: MIT Press.
- Gärdenfors, P., & Williams, M.-A. (2001). *Reasoning about categories in conceptual spaces*. In proceedings of the international joint conference of artificial intelligence (pp. 385–392).
- Gärdenfors, P., and F. Zenker. 2013. Theory change as dimensional change: Conceptual spaces applied to the dynamics of empirical theories. *Synthese* 190: 1039–1058.
- Gower, J.C., and G.B. Dijksterhuis. 2004. *Procrustes problems*. New York, NY: Oxford University Press.
- Indow, T. 1988. Multidimensional studies of Munsell color solid. *Psychological Review* 95: 456–470.
- Kant, I. 1785. *Grundlegung zur metaphysik der sitten*. Riga, Latvia: Hartknoch.
- Klahr, D. 1969. A Monte-Carlo investigation of the statistical significance of Kruskal's nonmetric scaling procedure. *Psychometrika* 34: 319–330.
- Kruskal, J.B., and M. Wish. 1978. *Multidimensional scaling*. Beverly Hills, CA: Sage.
- Lee, M.D., and K.J. Pope. 2003. Avoiding the dangers of averaging across subjects when using multidimensional scaling. *Journal of Mathematical Psychology* 47: 32–46.
- Lokhorst, G.J.C. 2018. Martin Peterson: The ethics of technology: A geometric analysis of five moral principles. *Science and Engineering Ethics* 24 (5): 1641–1643.
- Mair, P., I. Borg, and T. Rusch. 2016. Goodness-of-fit assessment in multidimensional scaling and unfolding. *Multivariate Behavioral Research* 51: 772–789.
- Mill, J. S. (1861/1965—Utilitarianism. In J. Robson (Ed.), *Collected Works of John Stuart Mill* (33 volumes). Toronto, Canada: University of Toronto Press.
- Miller, G.A. 1969. A psychological method to investigate verbal concepts. *Journal of Mathematical Psychology* 6: 169–191.
- Peterson, M. 2017. *The ethics of technology: A geometric analysis of five moral principles*. Oxford, UK: Oxford University Press.
- Peterson, M. 2018. The ethics of technology: Response to critics. *Science and Engineering Ethics* 24: 1645–1652.
- Core Team, R. 2017. *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing Retrieved from <https://www.R-project.org/>.
- Regier, T., P. Kay, and R.S. Cook. 2005. Focal colors are universal after all. *Proceedings of the National Academy of Sciences* 102: 8386–8391.
- Ross, W.D. 1930. *The right and the good*. Oxford, UK: Clarendon Press.

- Shepard, R.N. 1964. Attention and the metric structure of the stimulus space. *Journal of Mathematical Psychology* 1: 54–87.
- Shepard, R.N., and P. Arabie. 1979. Additive clustering: Representation of similarities as combinations of discrete overlapping properties. *Psychological Review* 86: 87–123.
- Shrader-Frechette, K. (2017). Review of Martin Peterson: The ethics of technology. *Notre Dame Philosophical Reviews*. University of Notre Dame, October 30, 2017. Web. November 11, 2017.
- Spearman, C. 1904. The proof and measurement of association between two things. *American Journal of Psychology* 15: 72–101.
- Spence, I., and J.C. Ogilvie. 1973. A table of expected stress values for random rankings in nonmetric multidimensional scaling. *Multivariate Behavioral Research* 8: 511–517.
- Stenson, H.H., and R.L. Knoll. 1969. Goodness of fit for random rankings in Kruskal's nonmetric scaling procedure. *Psychological Bulletin* 71: 122–126.
- Storms, G., P. De Boeck, and W. Ruts. 2000. Prototype and exemplar based information in natural language categories. *Journal of Memory and Language* 42: 51–73.
- Sturidsson, K., N. Långström, M. Grann, G. Sjöstedt, U. Åsgård, and E.M. Aghede. 2006. Using multidimensional scaling for the analysis of sexual offence behaviour: A replication and some cautionary notes. *Psychology, Crime & Law* 12: 221–230.
- Sturrock, K., and J. Rocha. 2000. A multidimensional scaling stress evaluation table. *Field Methods* 12: 49–60.
- Takane, Y., F. Young, and J. De Leeuw. 1977. Nonmetric individual differences multidimensional scaling: An alternating least squares method with optimal scaling features. *Psychometrika* 42: 7–67.
- Valentine, T., M.B. Lewis, and P.J. Hills. 2016. Face-space: A unifying concept in face recognition research. *Quarterly Journal of Experimental Psychology* 69: 1996–2019.
- Verbeemen, T., W. Vanpaemel, S. Pattyn, G. Storms, and T. Verguts. 2007. Beyond exemplars and prototypes as memory representations of natural concepts: A clustering approach. *Journal of Memory and Language* 56: 537–554.
- Verheyen, S., and P. Égré. 2018. Typicality and graded membership in dimensional adjectives. *Cognitive Science* 42: 2250–2286.
- Voorspoels, W., G. Storms, J. Longenecker, S. Verheyen, D.R. Weinberger, and B. Elvevåg. 2014. Deriving semantic structure from category fluency: Clustering techniques and their pitfalls. *Cortex* 55: 130–147.
- White, A., W. Voorspoels, G. Storms, and S. Verheyen. 2014. Problems of reliability and validity with similarity derived from category fluency. *Psychiatry Research* 220: 1125–1130.
- Zenker, F., and P. Gärdenfors. 2015. *Applications of conceptual spaces: The case for geometric knowledge representation*. Springer International Publishing Switzerland.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Steven Verheyen^{1,2} · Martin Peterson³

¹ Faculty of Psychology and Educational Sciences, KU Leuven, Leuven, Belgium

² Department of Psychology, Education and Child Studies, Erasmus University Rotterdam, Rotterdam, the Netherlands

³ Department of Philosophy, Texas A&M University, College Station, TX, USA