

Title page

A Cross-Cultural Assessment of the Semantic Dimensions of Intellectual Humility

Markus Christen^{a,*}, Mark Alfano^b, Brian Robinson^c

^a University Research Priority Program Ethics, University of Zurich, Zollikerstrasse 117, 8008 Zurich, Switzerland; christen@ethik.uzh.ch

^b Faculty of Ethics & Philosophy of Technology, Delft University of Technology, Jaffalaan 5, 2628 BX Delft, The Netherlands; mark.alfano@gmail.com

^c Philosophy Department, Michigan State University, East Lansing MI 48824, USA; prof.brobinson@gmail.com

* Corresponding author: Phone: +41 44 634 83 73 / Fax: +41 44 634 85 07

Manuscript

A Cross-Cultural Assessment of the Semantic Dimensions of Intellectual Humility

Abstract: *Intellectual humility can be broadly construed as being conscious of the limits of one's existing knowledge and capable to acquire more knowledge, which makes it a key virtue of the information age. However, the claim "I am (intellectually) humble" seems paradoxical in that someone who has the disposition in question would not typically volunteer it. There is an explanatory gap between the meaning of the sentence and the meaning the speaker expresses by uttering it. We therefore suggest analyzing intellectual humility semantically, using a psycholexical approach that focuses on both synonyms and antonyms of 'intellectual humility'. We present a thesaurus-based methodology to map the semantic space of intellectual humility and the vices it opposes as a heuristic to support philosophical and psychological analysis of this disposition. We performed the mapping both in English and German in order to test for possible cultural differences in the understanding of intellectual humility. In both languages, we find basically the same three semantic dimensions of intellectual humility (sensitivity, discreetness, and knowledge dimensions) as well as three dimensions of its related vices (self-overrating, other-underrating and dogmatism dimensions). The resulting semantic clusters were validated in an empirical study with English (n=276) and German (n=406) participants. We find medium to high correlations (0.54-0.72) between thesaurus similarity and perceived similarity, and we can validate the labels of the three dimensions identified in the study. But we also find indications of the limitations of the thesaurus methodology in terms of cluster plausibility. We conclude by discussing the importance of these findings for constructing psychometric scales for intellectual humility.*

Keywords: Intellectual humility; psycholexical analysis; semantics; synonymy; antonymy; Thesaurus databases

1. Introduction

Environmental ethicists interested in virtues have argued that, in an era of increasing interdependence, characterized by inscrutable and dynamic complexity, new virtues may be called for (Jamieson 2007; Fairbanks 2010; Hill 1999). Like the project of dealing with climate change at a global level, building, managing, maintaining, and simply engaging with Internet technology may require such novel virtues. Comparable to the global climate, the Internet is a hugely complex and dynamic system – one that we cannot simply study but with which we also interact and shape. Chief among the virtues championed by environmental virtue theorists is humility, which helps us to recognize and cope with our staggering limitations when dealing with complex systems. We believe that intellectual humility (IH) is an appropriate

counterpart in the context of information technology. In this paper, we present two studies that help to clarify the nature of intellectual humility in a cross-cultural setting.

The study of personality and conceptions of personality has been pursued by psychologists and other researchers in various ways, including observations in laboratory settings and field experiments, correlational studies of survey responses, and psycholexical analyses. The present research embodies the latter two methodologies and is informed by both philosophical theory and mathematical modeling tools developed in the physical sciences.

Psycholexical analysis dates back to Francis Galton's *Measurement of Character* (1884). The basic idea is that, all else being equal, a natural language is more likely to include a predicate for a property to the extent that the property is important to those who speak the language. English has the word 'defenestration' because it's been important to be able to talk about events in which someone is thrown out a window. English lacks a word for someone being thrown *in* a window because, over the last several centuries, such events didn't seem sufficiently worth talking about. This is not to say that every phrase or term refers. There are no unicorns despite the existence of the term 'unicorn'. Nor is it to say that everything worth talking about is already represented by a phrase or singular term. Words are sometimes coined because new phenomena come into existence or become important; terms are also sometimes coined because extant phenomena could not otherwise be parsimoniously described and explained (e.g., 'sexual harassment' – see Fricker 2007). In addition, sometimes a speaker coins words to describe or explain phenomena for which a word already exists, but of which the coiner is ignorant. In this way, words that are synonyms (or nearly so) emerge, further emphasizing the importance of the phenomena referred to. Regardless, the rough generalization that there is a strong positive correlation between the importance of phenomena in the lives of the speakers of a language and the probability of the existence of a term in the language that refers to those phenomena is hard to deny. Despite shortcomings with the psycholexical approach, there is broad agreement that it would be foolhardy "to ignore such a storehouse of accumulated wisdom as a natural starting-point for the study of behavioral attributes" (Wiggins 1937, p. 329). If this is on the right track, studying psychological language is an indirect way of studying the psychological properties people care about.

Psychologists in the psycholexical tradition don't stop there, though. They also typically argue that the *semantic structure* of a language reflects to some extent the *perceived structure of the phenomena* described by the language. In personality psychology, this insight was famously used by Allport & Odbert (1936) to create a semantic taxonomy of thousands of personality-relevant terms, which they argued represents how people conceive of personality. Now, the step from language to people's conception of personality is not identical to the step from their conception of personality to *actual* personality, but it's natural to think that there will be at least a positive correlation – if only a weak one – between how people think about personality and how personality actually is. This two-step connection (from language about personality to conceptions of personality, from conceptions of personality to actual personality) has been empirically validated by personality models such as the Big Five (Peabody & Goldberg 1989) and Big Six (Ashton et al. 2004; Saucier 1997). Existing archival work in this tradition has traditionally relied on the *dictionary* as a source. Our approach in this paper is

unusual in relying on the *thesaurus* as a source.¹ This methodological innovation allows us to explore new data with new methods.

The Big Six includes an H factor that represents facets of personality broadly related to honesty and humility. *Intellectual* humility seems to involve a consciousness of the limits of one's knowledge, including sensitivity to circumstances in which one's native egocentrism is likely to function self-deceptively (Roberts & Woods 2007), though others regard it as more of a "second-order" open-mindedness (Spiegel 2012). In our age of information, intellectual humility has grown all the more relevant. However, little conceptual or empirical work has explored this trait. We think that the psycholexical approach is especially promising in the investigation of intellectual humility because self-report questionnaires are likely to be especially unreliable as measures of this construct. Someone who is genuinely (intellectually) humble is unlikely to report being (intellectually) humble, and someone who reports being (intellectually) humble is unlikely to be (intellectually) humble. Humility – whether intellectual, moral, or otherwise – seems to involve a paradox of self-attribution.

Additionally, our investigation is motivated by Aristotle's insight, reiterated in contemporary philosophy by Roberts & Wood (2007), that a virtue (which, in this paper, we construe simply as a positive value-laden personality disposition or dimension of individual difference) is often best understood in the context of related virtues and the vices they oppose. Put a different way, by contextualizing a term for a virtue in the constellation of its near-synonyms and its near-antonyms, we can generate a perspicuous representation of the meaning of the term.

We furthermore are interested in potential cultural differences regarding the understanding of intellectual humility using the psycholexical approach. The first successful psycholexical replication of the Big Five model outside of English was in German (Ostendorf 1990). This suggests that our thesaurus-based methodology is especially likely to replicate in German, though we hope eventually to study other languages, including non-Germanic and even non-Indo-European languages.

For these reasons, we investigated the trait of intellectual humility psycholexically by comparing 'intellectual humility' (German: 'intellektuelle Redlichkeit' or 'intellektuelle Bescheidenheit', see remarks in Section 2.1) with both its antonyms and synonyms in two languages: English and German. In Section 2, we will first outline our methodology, which relies on thesaurus databases to create similarity metrics among terms, and which is validated in a survey study. In Section 3, we describe the theoretical basis of an algorithm that we use for visualizing the semantic space of intellectual humility. Section 4 shows the results of both the thesaurus analysis and the validation study. In Section 5, we discuss the relevance of our findings for studying intellectual humility in a cross-cultural context and point to limitations of our approach.

¹ Interestingly, Galton (1884) used Roget's *Thesaurus* as a source, but since then almost all psycholexical research has used the dictionary.

2. Material and Method

2.1. Thesaurus analysis

Our analysis is based on the assumption that the practice of speaking a language is precipitated in dictionaries, lexicons, and other wordbooks. Of particular interest is the thesaurus – a language reference book or database organized to help its users find words related to a concept but having slightly different shades of meaning or connotation. Thesaurus dictionaries have a long tradition, starting in the 17th century and cumulating in famous books like Rogge’s *Thesaurus*, published in 1852, an outstanding work of English lexicography (Hüllen, 2004). Thesauruses reflect what people in their daily use of language – in particular when writing – consider semantically similar to a given term. They can be understood as expressions of “practical synonymy”, which involves employing the principle of synonymy for semanticizing lexemes, i.e. basic units of lexical meaning that exists regardless of the number of inflectional endings they may have or the number of words they may contain (Hüllen, 2004).

There is a rich theoretical tradition regarding the notion of synonymy in linguistics, philosophy of language, and other fields. On a strict understanding, synonymy refers to the fact that there may be several different words for expressing exactly the same meaning, which is an understanding that is difficult to uphold, as Quine (1951) has observed. Within the field of semantics, various notions of synonymy like semantic fields (Trier, 1931), the structuralist investigations of Harris (1973), and the pragmatic suggestion of Jones (1986) have been developed. In addition, databases such as *WordNet* (see <http://wordnet.princeton.edu/>), which labels the semantic relations among words, have been developed.

Today’s thesaurus databases list synonyms in a broad sense, i.e., they employ some notion of “meaning similarity”. The major aim of a thesaurus is not to find a replacement Y of a certain term X what has exactly the same meaning of X, but rather to find a term Y that has a slightly different shade of meaning for better expressing what the writer actually wants to express. Thus, a thesaurus is broader in capturing word relationships beyond synonymy in a strict sense, while still being more specific than the mere co-occurrence statistics of terms in texts. Contemporary thesauri, then, constitute an ideal corpus for the Ramsification method we propose, use, and assess from an ethical point of view in a series of recent papers (Alfano 2015; Christen, Robinson & Alfano 2014; Christen, Alfano, Bangerter, & Lapsley 2013). The basic idea behind this method is to make holistic inductive inferences about the meaning and structure of concepts from patterns of platitudinous truths and inferential connections (Lewis 1966, 1970, 1972; Ramsey 1931).

The present research explores the semantic space of intellectual humility by first identifying the most common synonyms and antonyms of ‘intellectual humility’. Next, by referring to the thesaurus.com database (the largest online thesaurus for American English), we associate each identified term with a *word-bag* $T = \{t, t_{syn1}, t_{syn2}, t_{syn3}, \dots, t_{syn n}\}$, which is the set of synonyms listed for that term including t itself. By comparing word bags, we then create a similarity metric by calculating the relative overlap of each pair of word-bags (see Section 3.1. for

technical details). The similarities calculated in this way are then used in a novel clustering and visualization tool that generates a semantic map of the terms involved.

More specifically, we used the following 7-step procedure (see also Fig. 1):

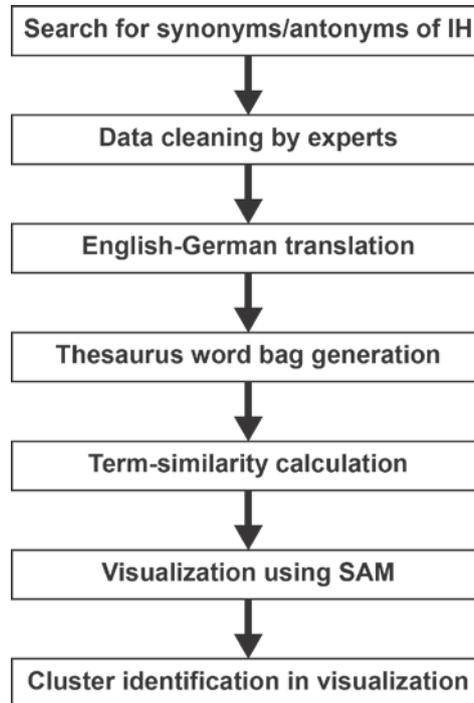


Figure 1: Outline of the procedure for generating the semantic maps of IH synonyms and antonyms

- 1) We identified potential synonyms and antonyms for ‘intellectual humility’ in three ways:
 - a. In summer 2013, we searched philosophy and psychology journals for articles that discuss intellectual humility; we found 24 papers or related texts (such as calls for proposals, abstracts, and papers).
 - b. Also in summer 2013, we performed an Internet search for entries on ‘intellectual humility’ and found 20 entries that dealt in a significant way with the concept.
 - c. We identified scales that are used in psychology for constructs that have some similarity to intellectual humility (e.g., the H factor of the Big Six personality inventory).

In all these texts, we identified terms that are used to represent the meaning of ‘intellectual humility’ or its relevant vices.

- 2) Four raters who have experience with the philosophical topic of intellectual humility assessed all terms collected in step 1 to determine whether they could be used to express the concept of intellectual humility or a related vice. A term was kept on the list if three out of four raters agreed to do so. In this way, we identified 52 synonyms and 69 antonyms for ‘intellectual humility’. Each term was represented at least in noun form and usually in adjective form as well: for example, {tolerance, tolerant}.

- 3) The notion of “intellectual humility” is uncommon in German, partly because the German translation of “humility” generates three terms with different meanings (‘Bescheidenheit’, ‘Demut’, ‘Ergebenheit’), which back-translate into a diverse English term set {abjection, devotedness, devotion, faithfulness, humbleness, humility, lowliness, loyalty, meekness, modesty, pudency, resignedness, subjectedness, submissiveness, unassumingness, uxoriousness}.² This phenomenon exemplifies the well-known indeterminacy of translation thesis of Quine (1960, chapter 2). The meaning of ‘intellectual humility’ can be captured by the German ‘intellektuelle Redlichkeit’ or ‘intellektuelle Bescheidenheit’, which would represent quite different facets of the English term. However, as our study does not rely on a single term (where the problem of translation ambiguity is inevitable), but on a whole constellation of terms, our approach can handle this issue. We have translated all English synonyms and antonyms in such a way as to compensate for ambiguity on a single-term level. For example, the English term ‘mindfulness’ translates into the German terms ‘Achtsamkeit’ and ‘Aufmerksamkeit’; and their back-translation reveals (among others) the English term ‘attentiveness’, which is another term identified in step 2. By choosing appropriate pairs out of these ambiguous sets we were able to include many of the possible translations of the English synonym and antonym sets of IH, which diminishes the problem of translation ambiguity and vagueness.
- 4) We identified all entries for each term generated in step 2 in the thesaurus.com database to generate word-bags for each synonym and antonym. For example, the word-bag for ‘tolerance’ included all entries on thesaurus.com for the term set {tolerance, tolerant}. We did the same for all German synonyms and antonyms, using the largest German thesaurus provider, Woxikon (<http://synonyme.woxikon.de/>). This step revealed that three German antonyms (‘Egozentriertheit’, ‘Prahlsuch’ and ‘Prunksucht’) had no entry in the thesaurus database, so those terms were deleted from the German antonym list.
- 5) Next, we calculated the similarity in overlap between every pairwise combination of word-bags in either language. The technical details of calculation are outlined in Section 3.1.
- 6) The similarity measures obtained in this way were then used as inputs in a visualization algorithm called *superparamagnetic agent mapping*, which employs self-organizing agents governed by the dynamics of a clustering algorithm inspired by spin physics to generate low-dimensional representations for which the characteristics of nonlinear data structures are preserved or even emphasized. The technical details of calculation are outlined in Section 3.2.
- 7) Finally, using the same clustering paradigm in an adapted version from (Ott et al. 2005), we identified clusters on the map generated in step 6. The terms of the clusters identified in this way were used in the survey study. In this way we were able to reduce the number of synonyms from 52 to 23 and the number of antonyms from 69

² We used the LEO online dictionary, one of the largest and most popular German-English dictionaries (<http://dict.leo.org>).

(English) or 66 (German) to 33. This was necessary to make the survey study a feasible length.

The final maps resulting from this procedure are used to inform our reasoning about intellectual humility.

2.2. Survey Study

Our approach is based on the assumption that the entries in thesaurus databases reflect similarities of terms based on their usage in (written) language in a rather broad language community, neglecting differences due, for example, to different dialects. At the outset, it was not clear whether the result of an analysis based on this assumption would lead to clusters that people would consider “similar” according to a pre-theoretical understanding of the terms.

We therefore performed two online survey studies – one in English (USA) and one in German (Switzerland) – designed to empirically test whether people would cluster synonyms and antonyms in a similar way as the groups we separately identified on the thesaurus maps. In the questionnaire, the participants first provided informed consent and basic demographic information (age and gender). Then they were presented with a test example to help them understand the task. They were given a list of 10 terms and three boxes. Their task was to drag the terms to the boxes using the mouse pointer such that all terms in a single box are similar in meaning according to the participants’ understanding. They were instructed that the groups may vary in size, that they don’t have to attribute all terms and that the order of the terms within a box is not relevant. They were allowed to regroup the terms (e.g., shifting a term from box 1 to box 2) if they believed a better sorting could be achieved in this way. Finally, participants were asked to name the groups and to rate how similar the members of a single group are to each other using a 6-point Likert scale (endpoints: the terms of one group are very similar / although being in the same group, the terms are still quite distinct). We also provided a possible solution of the sorting in the test example to ensure that the participants understood the task. After the test example, the participants received 23 IH synonyms identified in step 7 (see Section 2.1.) and an additional distractor term (‘ugliness’, ‘Hässlichkeit’) for sorting into three groups. Next, they received 33 IH antonyms and an additional distractor term (‘beauty’, ‘Schönheit’) for sorting into five groups. (The number of groups is based on the number of clusters for synonyms and antonyms, i.e., 3 and 5 respectively we found separately in the thesaurus analysis, as explained in Section 4.1.) We included distractor terms to check for sorting accuracy. Our hypothesis was that the distractor term would usually not be attributed and thus would remain on the original list (i.e., not placed in one of the boxes).

The survey among participants for the English survey was conducted in January 2015 using Amazon Mechanical Turk (restricted to participants with a United States IP address). In this sampling method, participants received a small reimbursement for completion of the survey. The survey for the German value groups was conducted in March 2015 using a service provided by the University of Zurich that addresses students, faculty, and staff of the university. Here, a lottery (book gift cards) was chosen as incentive for participation. The study was cleared in accordance with the ethical review processes of both universities.

3. Theory and Calculation

3.1. Similarity Measure

An important heuristic in the method outlined above consisted of a visualization tool to help the experts in evaluating the clusters. The idea was to map the similarities between the meanings of the terms represented by word bags on a two dimensional plane such that similarities translate into distances on the map. Therefore, a similarity measure was needed as input for the visualization tool.

For comparing word bags of the terms, the similarity measure was obtained by calculating the relative overlap of each pair of word bags. Let $T = \{t, t_{syn1}, t_{syn2}, t_{syn3}, \dots, t_{syn n}\}$ be the word bag of a term t . Then the similarity S of two terms t^1 and t^2 is defined as:

$$S(t^1, t^2) = \frac{|T^1 \cap T^2|}{\min\{|T^1|, |T^2|\}}$$

The similarity of two terms equals 1 when the word bag of one term overlaps completely with the word bag of another term. The similarity is 0 if the word bags are mutually distinct. For example, the word-bag of ‘tolerance’ contains 55 terms and the word-bag of ‘broadmindedness’ contains 40 terms. 12 terms are contained in both word-bags. Hence, the similarity between ‘tolerance’ and ‘broadmindedness’ is $12/40 = 0.3$.

Although S is a similarity and not a distance measure,³ it is adequate for our visualization method (see Ott et al. 2014 for details). This allowed us to generate a dissimilarity matrix (1 minus the similarity value) that served as input for the superparamagnetic agent map (SAM) algorithm (see below).

To empirically validate the similarity maps obtained from thesaurus data, we needed a measure of similarity for participant-generated data. Here, we calculated the ratio of how often two terms were classed together in the same group by the participants compared to the total number of participants. For example if 79 out of 276 English-speaking participants placed ‘attentiveness’ and ‘circumspection’ in the same group, so the similarity between the two terms is $79/276 \approx 0.29$. Although the similarity measures obtained in this way are not expected to match the ones of thesaurus similarity (e.g., it is much less likely to obtain a similarity value of 0 in the survey), we hypothesized that the measures would correlate if the participant ratings reflect thesaurus similarity.

3.2. Superparamagnetic Agent Map

Visualizing a semantic space is a standard problem of dimensionality reduction. Classical

³ A distance measure has to fulfill the triangle inequality, i.e. $d(a, b) + d(b, c) \geq d(a, c)$, where $d(a, b)$ stands for the distance between points a and b . This inequality is violated by our measure $S(t^1, t^2)$. In the following, we use the term “dissimilarity” to denote $1 - S(t^1, t^2)$ in the original space, whereas on the map, we use the term “distance”, as we refer to the Euclidean distance of points on the two-dimensional plane.

approaches aim to represent the data structure on a linear subspace of the original data space. For example, principal component analysis performs a projection onto the axes with maximal data variance, whereas the goal of multidimensional scaling is to find a low-dimensional embedding that preserves the inter-point distances. These methods often perform poorly when applied to nonlinear data structures. Furthermore, for many real-world applications, data vectors are not available. Instead, researchers are faced with similarity or proximity data, as in our case. We therefore have applied a novel visualization tool called Superparamagnetic Agent Mapping (SAM) (Ott et al. 2014).

To conceptualize this mapping, imagine each term as a particle that naturally repels all other particles. However, as overlap between two terms increases, they become more attracted to each other. Thus, SAM typically produces clustering, where several particles clump together (connoting similarity) while collectively repelling a different cluster (connoting collective difference between the two clusters).

More formally, the method is an iterative two-step procedure that is repeated until a threshold condition is reached. In the first step, each data item (= synonym or antonym term) is represented by a Potts spin variable and the dissimilarity matrix is encoded in the spin couplings. The spin system is treated in the formalism of the canonical ensemble, giving the probability for a certain spin configuration. One can then observe that the spins whose corresponding data items are similar tend to cluster in terms of the pair correlation G_{ij} , i.e., the probability of two spins being in the same state. By introducing a temperature-like parameter T , a cluster hierarchy can be generated. For smaller T , all spins tend to be in the same state. Upon an increase in T , large clusters break up into smaller clusters in a cascade of (pseudo-)phase transitions

In the second step, each data item is represented by an agent in a two-dimensional coordinate system. The agents move according to laws that are governed by the local interactions of the spin system calculated in step 1. In order to calculate G_{ij} , a Markov chain Monte Carlo algorithm is employed, which generates a sequence of binary pair correlation states $G_{ij}(t) \in \{0,1\}$. Starting from a random distribution, two agents move towards each other if $G_{ij} = 1$, i.e., if the corresponding spins are in the same state in the current configuration, otherwise the agents drift apart, leading to a two-dimensional distribution of agents. For the precise formalism, we refer to Ott et al. (2014). It has been shown (Ott et al. 2014) that SAM is superior to standard methods such as factor analysis, principal components analysis, and multidimensional scaling in preserving the topology of the data space with clustered data.

To quantitatively determine which terms on the map actually form a group, we calculated the Euclidean distance of each data item on the map and used the resulting distance matrix for the sequential superparamagnetic clustering algorithm (which is basically equivalent to step 1 of the SAM procedure). This clustering algorithm has several theoretical advantages. In particular, by extracting the most stable sets and re-clustering, the algorithm overcomes the problem that natural entities or classes that are heterogeneously represented in the feature space have to be selected against their respective background. Some classes therefore are not found in the original set, but only in a subset of adequate homogeneity.

4. Results

In the subsections below, we showcase the results of both the thesaurus study and the survey study.

4.1. Superparamagnetic Agent Maps of IH synonyms and antonyms

We produced three maps to convey our results. Figure 2 is the synonym map, showing the degree of overlap among intellectual humility’s 52 synonyms. The terms predominantly cluster into three groups. The first group (displayed in green) outlines a *Sensibility dimension* and is exemplified by terms such as ‘attentiveness’, ‘comprehension’ and ‘mindfulness’. We take this cluster to be representative of the notion that an intellectually humble person will be open and responsive to new ideas and information. The second (red) cluster we call the *Knowledge dimension*; it is illustrated by terms such as ‘curiosity’, ‘research’, and ‘thirst of knowledge’. The difference between these two dimensions indicates that there is an intuitive difference between seeking new information or ideas (knowledge dimension) and being open to them when they are presented (sensibility dimension), since new information sometimes appears without being sought. Third, we have named the blue cluster the *Unpretentiousness dimension*, which is typified by ‘humility’, ‘unostentatiousness’, and ‘unpretentiousness’. Finally, some terms (shown in black) have intermediate positions among these groups (e.g., ‘flexibility’ and ‘tolerance’) and do not fit neatly within any cluster.

Similar dimensions emerged in the German data set, although the content of the clusters (in terms of translations) differs. For example, the unpretentiousness dimension contains 9 terms in English and 10 terms in German. Five terms match in either language (e.g., ‘integrity’ and ‘Integrität’) – although we have to bear in mind that “matching” should be taken with a grain of salt due to indeterminacy in translation. A noteworthy semantic difference is that in the German map, the unpretentiousness dimension also includes terms like ‘flexibility’ and ‘openness’ that are not in the English cluster, whereas the English cluster includes terms like ‘decency’ and ‘demureness’, whose translations (‘Anstand’ and ‘Zurückhaltung’) are located in the German sensibility dimension. Most interesting is the lack of overlap in the knowledge dimension: the English cluster contains terms that refer to the generation of knowledge, whereas the German cluster contains terms that refer to having or using knowledge – the translation of those terms show up in the sensibility dimension in the English map.

Figure 3 shows the results of the antonym map, displaying the degree of overlap between intellectual humility’s 69 (respectively 66 in German) opposites. The first result to notice is that in the English map, almost all the terms are aligned and cluster at each endpoint. We take this to represent the distinction between underrating and overrating. Clustering the map reveals five distinct groups along three dimensions. Two groups fall on the overrating dimension. The larger, orange cluster includes terms such as ‘vanity’, ‘pride’, and ‘arrogance’. This cluster suggests that one way not to be intellectually humble is to be overly focused on one’s own high status. The frozen limit, as it were, of this dimension consists of the red cluster that contains terms like ‘egoism’ and ‘self-centeredness’, which can be understood as pure self-centricity.

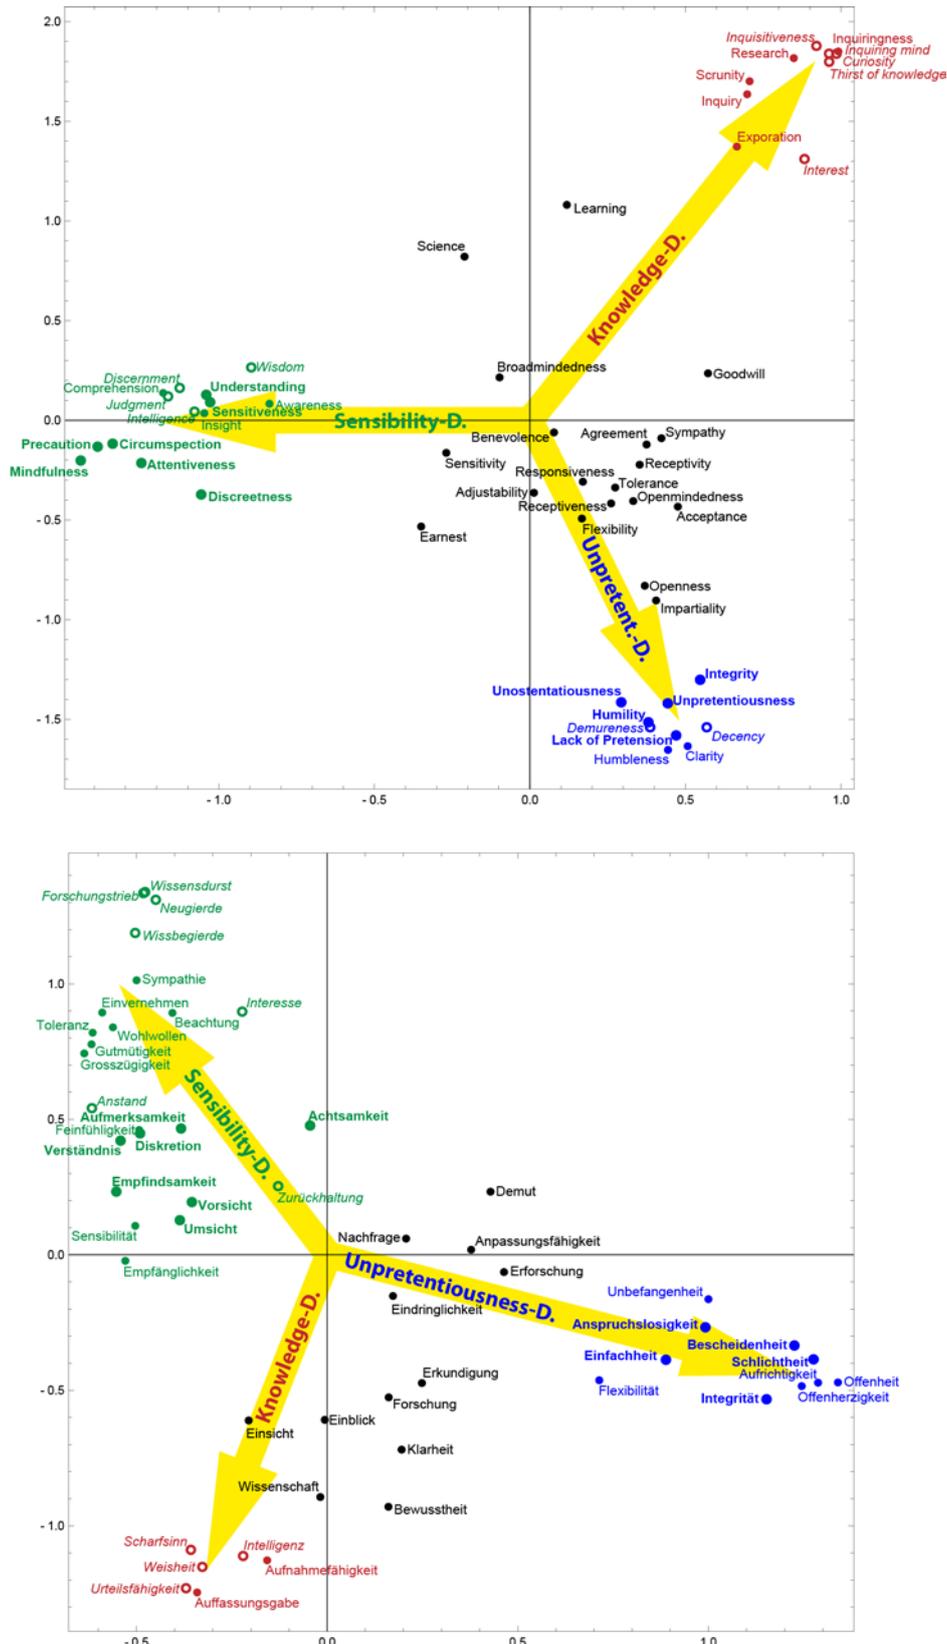


Figure 2: The IH Synonym maps in English and German show clusters, which are identified on the map are color-coded. Cluster terms printed in bold indicate translation pairs that are present in the same dimensions in either language. Ring-shaped points and terms printed in italic refer to terms whose translations are present in different dimensions in the other map.

OVERRATING oneself is not, however, the only way to fail to be intellectually humble. The opposite endpoint has two clusters that point into the underrating dimension. One cluster (dark blue), typified by terms such as ‘bias’, ‘prejudice’, and ‘unfairness’ is about underrating the other, whereas the light blue cluster, characterized by terms such as ‘diffidence’, ‘servility’, and ‘timidity’, is about underrating the self. These two groups are conceptually unified in that they both involve underrating someone. The key difference, however, is who is underrated. Underrating others is one way fail to be intellectual humble. Such underrating often goes with overrating oneself, though our map suggests that this needn’t be so since they are conceptually distinct. Alternatively, one can underrate oneself, suggesting that there is such a thing as being too humble, such that one’s lack of pride ceases to have any positive value and may even take on negative value. Finally, the clustering algorithm identified a fifth group that can be considered a separate dogmatism dimension, characterized by terms like ‘bullheadedness’ and ‘stubbornness’.

Again, the German map reveals similar dimensions, but also some notable differences to the English map when comparing the identified clusters. First, the dogmatism dimension singles out much more compared to the English map, where the terms are close to those of the “underrating others” dimension (which makes a certain amount of intuitive sense). Second, the two sub-groups of the overrating dimension seem to be less strongly related to each other compared to the English map, although the terms included in the clusters show a large degree of overlap (13 out of 18 terms are present in the clusters of both languages). Thus, most notably, the German map does not reveal a difference between the two aspects of underrating (self vs. other), which is present in the English map. Two (‘Schüchternheit’ and ‘Zaghaftigkeit’) out of six terms of the English sub-group show up in the German underrating cluster, whereas the other four are not present in any cluster.

Finally, we mapped all synonyms and antonyms together. We preserved the colors from the two previous maps, but we used a different typesetting to distinguish synonyms (Times New Roman, a serif font) from antonyms (Arial, a sans serif font). The resulting map retains many of the structural features of the previous maps, but with a few significant changes. In both maps, we find that the antonyms cluster in a more pronounced way into different groups representing the dimensions identified previously. This suggests that the antonym dimensions in general are more distinct compared to the synonym dimensions. In the English map, the two major dimensions (overrating vs. underrating) of the IH vices clearly show up as separate clusters, whereas the dogmatism dimension is closely related to the underrating dimension (but nevertheless still separated). Most interesting is that the sub-group “underrating self” is included in the part of the map that displays many IH synonyms, notably terms related to the unpretentiousness dimension. Within this part of the map the three dimensions of IH synonymy still show up at either edge of the large term cloud. The German map basically displays the same qualitative picture; the main difference is that the overrating subgroups (self vs. others) are at very different places in the map, indicating a larger difference between these two aspects of IH vices compared to the English map. Furthermore, terms that are (in the English map) related to the “underrating self” cluster that is not present in the German antonym map are also embedded in terms related to IH synonyms, but they are less clearly related to the unpretentiousness dimension.

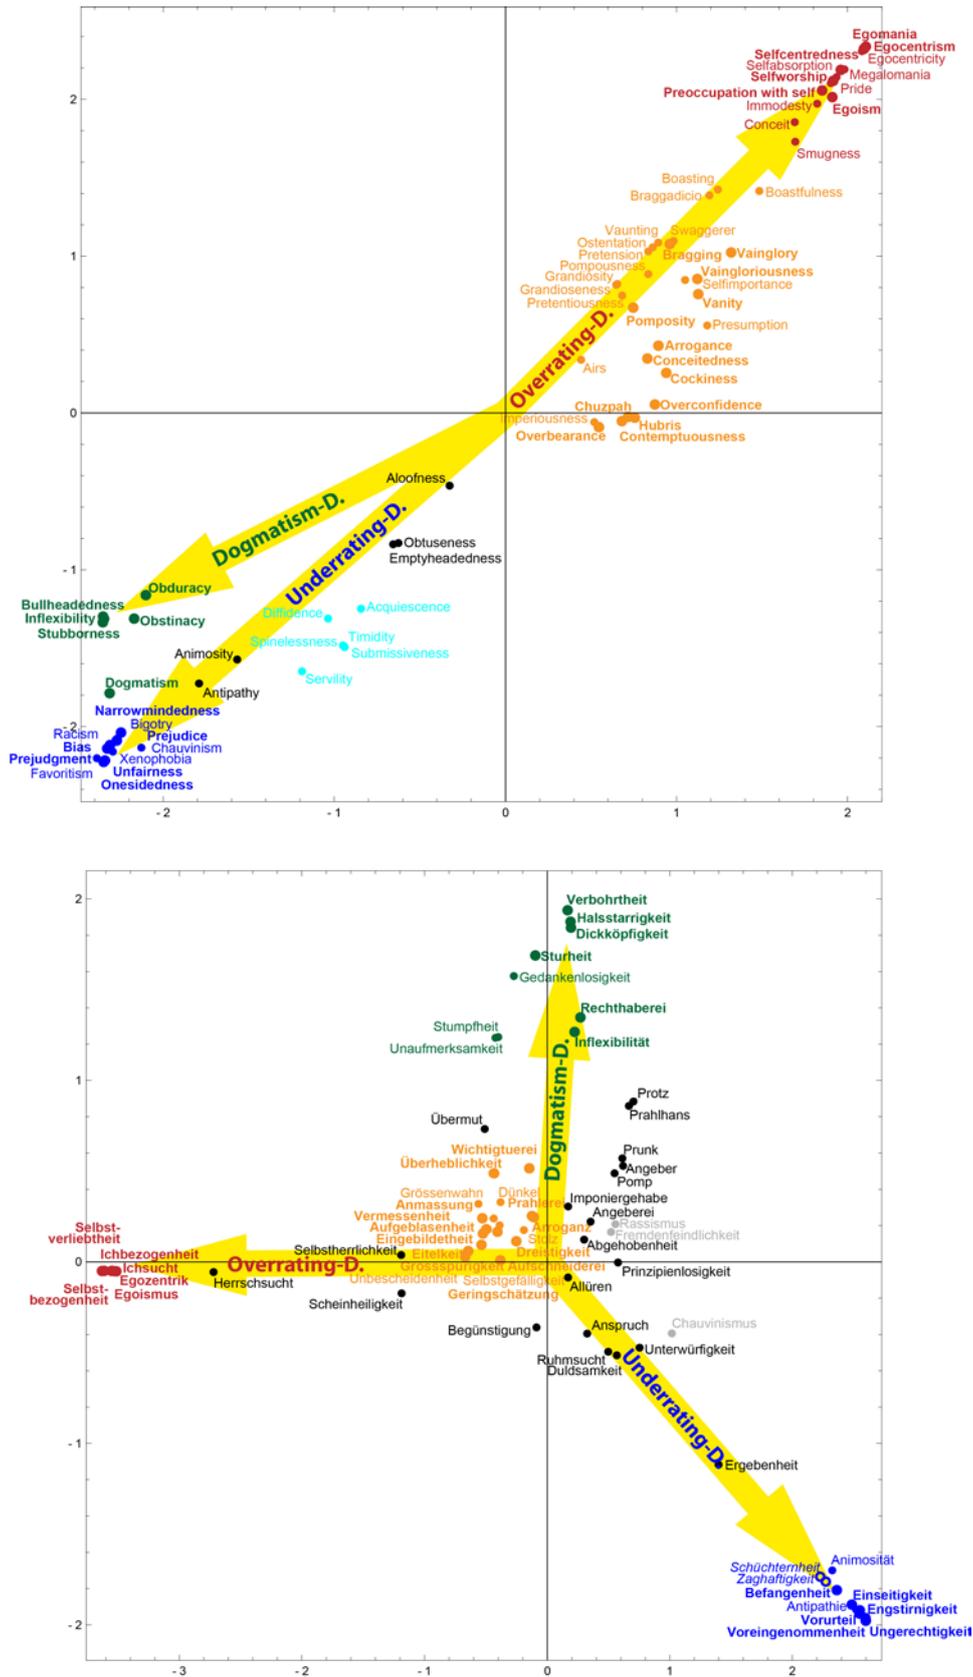


Figure 3: The IH Antonym maps use the same coloring and typographic conventions as in Figure 2. The grey terms ‘Rassismus’, ‘Fremdenfeindlichkeit’ and ‘Chauvinismus’ are special, because their word-bags do not overlap with those of any other term of the map, i.e. their placement on the map relative to the other terms is random and does not bear information.

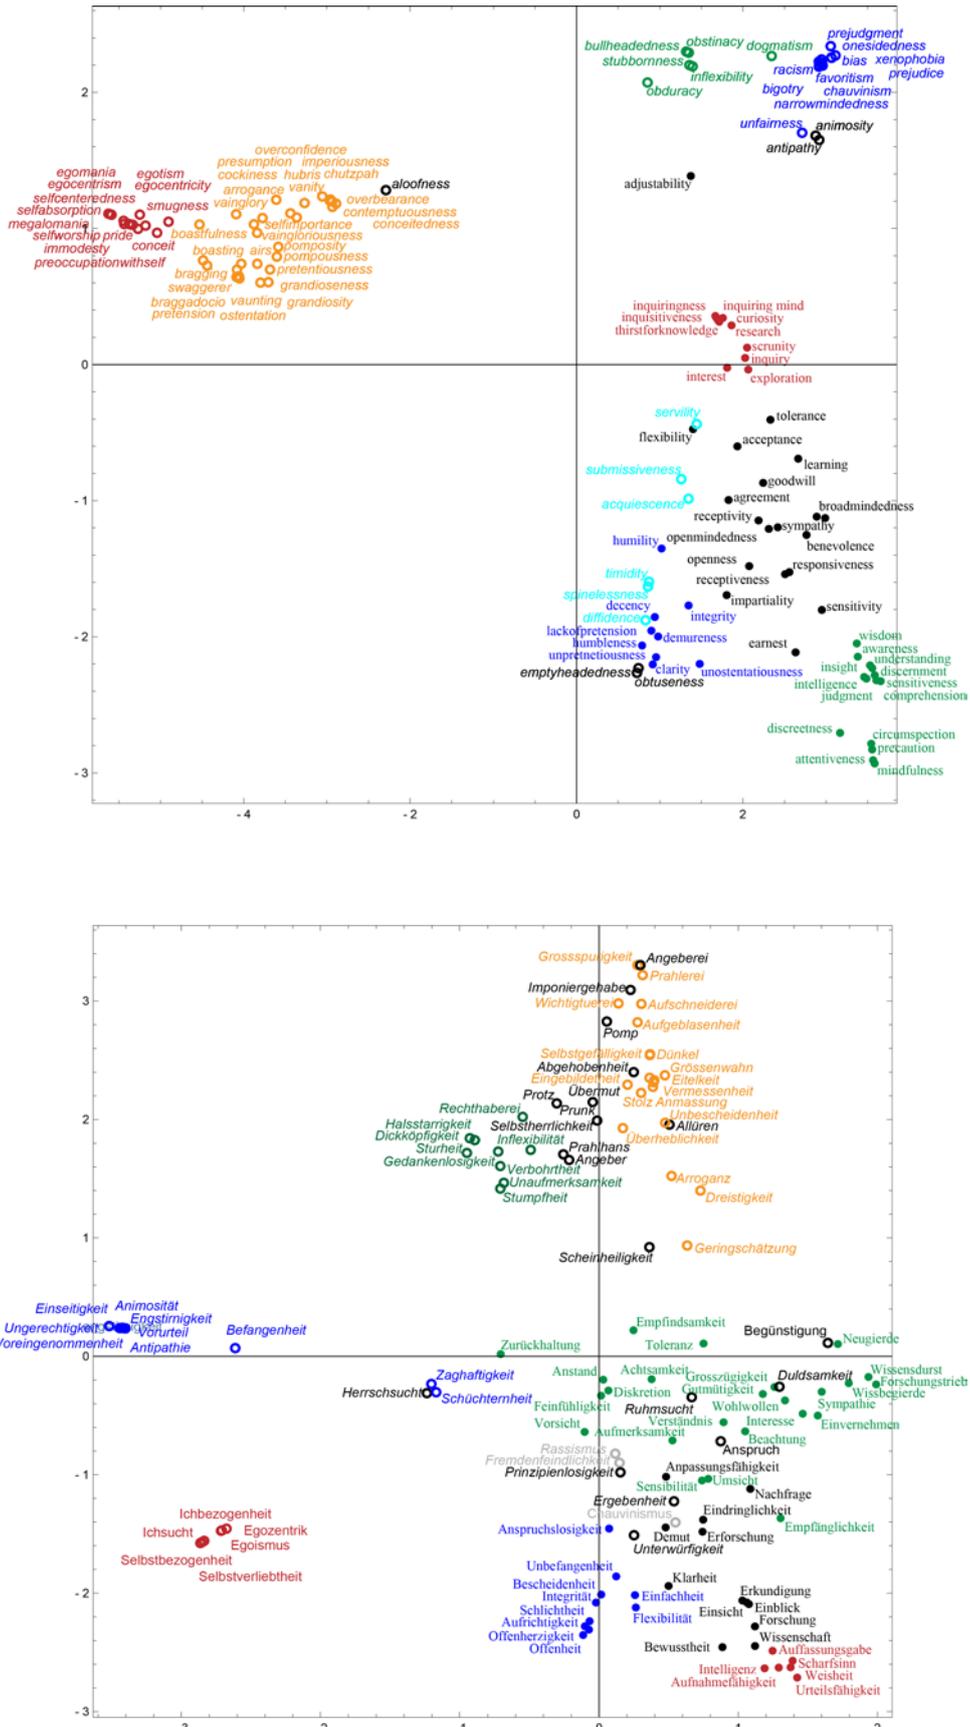


Figure 4: The maps displaying both IH synonyms and antonym map also follow the same color-coding conventions hold as in the other two maps. Synonyms are in Times New Roman (a serif-font) and full circle, antonyms are in Arial (a san-serif font) italic and open circles.

4.2. Survey Study

The survey study aimed to validate the thesaurus similarities and the clusters found. An interesting general result was that there was huge variability in how the people grouped the terms (see below). In particular, participants often built groups that included the distractor term. In English, 44.6% of the participants created groups that also included ‘ugliness’ (German: only 16.0% attributed ‘Hässlichkeit’) in the case of IH synonyms, and 68.8% included ‘beauty’ (often together with, e.g., ‘vanity’) in the case of IH antonyms (German: 31.5% attributed ‘Schönheit’). We interpreted this as meaning that our distractors were sub-optimal, so we abstained from taking allocations of the distractor as a “hard” criterion for excluding participants. Rather, we calculated the correlations for three different ways of taking the distractors into account:

- 1) In the first group, we included all participants who provided complete data sets (English: $n = 276$; German: $n = 406$)
- 2) For the second group, we checked in detail to which other term people usually attributed the distractor terms, and we identified those terms to which fewer than 5% of the participants made attributions. We then excluded all participants who made such “wrong” attributions (remaining: English: $n = 201$; German: $n = 326$)
- 3) The third group consists of only those participants who never allocated a distractor term to a group (remaining: English: $n = 72$; German: $n = 264$);

Using the data of the three groups, we calculated the pairwise similarity of all synonyms and antonyms as explained in Section 3.1, and we determined the Pearson correlation between the obtained values and the thesaurus similarity of each pair. The results are displayed in Table 1, which indicates that higher sorting accuracy (i.e. when excluding participants that did attribute the distractors) increases the correlation.

		Group 1	Group 2	Group 3
English	Synonyms	0.700	0.709	0.716
	Antonyms	0.694	0.698	0.707
German	Synonyms	0.473	0.486	0.541
	Antonyms	0.649	0.659	0.686

Table 1: Pearson correlation of thesaurus similarity and survey similarity ($p < 0.001$ in all cases)

To validate the clusters, we used the following approach in order to deal with the high variability of term grouping: we determined all different attribution of terms (in the following, we use the notion ‘pattern’ for a distinct grouping of terms by participants) the people made when allocating the terms to 3 (synonyms) or 5 (antonyms) groups. We only kept those patterns that were chosen by at least ~1% of the participants (in English: by 2 or more participants, in German: by 4 or more) and that included at least 3 elements for further analysis (see Table 2). In this way we reduced the large variability of term attribution to those patterns where people had some minimal agreement. Interestingly, the English participant showed considerably higher variability in generating patterns compared to the German participants.

		Total # of different patterns	# of different patterns chosen by at least ~1%	# frequent patterns divided by # all patterns
English	Synonyms	666	62	24.3%
	Antonyms	965	70	18.5%
German	Synonyms	648	43	41.9%
	Antonyms	962	56	37.4%

Table 2: Overview of variability and frequency of the term attribution patterns

In the frequent patterns, we calculated the similarity of each pair of terms by determining how often they are present in the same pattern. The resulting similarity matrix was clustered using a standard k -means clustering algorithm, as the number of clusters was predetermined ($k = 3$ in the case of synonyms and $k = 5$ in the case of antonyms; $k = 4$ for the German antonyms as well). The results are displayed in Tables 3 (English) and 4 (German). In general, the result reveals that the survey clusters are of greater semantic homogeneity compared to the thesaurus clusters: for example, the sensibility cluster in both languages is more focused on issues like mindfulness, and the knowledge-related terms are regrouped in the knowledge dimensions (there are exceptions, e.g. that ‘overbearance’ shows up in the underrating-self group). Interestingly, if one looks only for 4 clusters in the German antonyms, the same merging is discernible as seen in the thesaurus analysis, i.e. a single “underrating-cluster” emerges.

IH-Synonyms English		
	Thesaurus Cluster Members	Survey Cluster Members
Cluster 1: Sensibility	attentiveness, <i>circumspection</i> , discernment, discreetness, intelligence, judgment, <i>mindfulness</i> , precaution, sensitiveness, <i>understanding</i> , wisdom	<i>circumspection</i> , <i>mindfulness</i> , <i>understanding</i>
Cluster 2: Unpretentiousness	<i>decency</i> , <i>demureness</i> , <i>humility</i> , <i>integrity</i> , <i>lack of pretension</i> , <i>unostentatiousness</i> , <i>unpretentiousness</i>	<i>decency</i> , <i>demureness</i> , discreetness, <i>humility</i> , <i>integrity</i> , <i>lack of pretension</i> , precaution, sensitiveness, <i>unostentatiousness</i> , <i>unpretentiousness</i>
Cluster 3: Knowledge	<i>curiosity</i> , <i>inquiring mind</i> , <i>inquisitiveness</i> , <i>interest</i> , <i>thirst for knowledge</i>	attentiveness, <i>curiosity</i> , discernment, <i>inquiring mind</i> , <i>inquisitiveness</i> , intelligence, <i>interest</i> , judgment, <i>thirst for knowledge</i> , wisdom
IH-Antonyms English		
Cluster 1: OVERRATING (1)	<i>egocentrism</i> , <i>egomania</i> , <i>egoism</i> , <i>preoccupation with self</i> , <i>self-centeredness</i> , <i>selfworship</i>	conceitedness, <i>egocentrism</i> , <i>egomania</i> , <i>egoism</i> , <i>preoccupation with self</i> , <i>self-centeredness</i> , <i>selfworship</i> , vaingloriousness, vainglory, vanity
Cluster 2: OVERRATING (2)	<i>arrogance</i> , <i>bragging</i> , <i>chutzpah</i> , <i>cockiness</i> , conceitedness, contemptuousness, <i>hubris</i> , <i>overbearance</i> , <i>overconfidence</i> , <i>pomposity</i> , vaingloriousness, vainglory, vanity	<i>arrogance</i> , <i>bragging</i> , <i>chutzpah</i> , <i>cockiness</i> , <i>hubris</i> , <i>overconfidence</i> , <i>pomposity</i>
Cluster 3: Underrating other	<i>bias</i> , <i>narrow-mindedness</i> , <i>onesidedness</i> , <i>prejudgment</i> , <i>prejudice</i> , <i>unfairness</i>	<i>bias</i> , contemptuousness, dogmatism, <i>narrow-mindedness</i> , <i>onesidedness</i> , <i>prejudgment</i> , <i>prejudice</i> , <i>unfairness</i>
Cluster 4: Underrating self	<i>diffidence</i> , <i>timidity</i>	<i>diffidence</i> , <i>overbearance</i> , <i>timidity</i>
Cluster 5: Dogmatism	<i>bullheadedness</i> , dogmatism, <i>inflexibility</i> , <i>obduracy</i> , <i>obstinacy</i> , <i>stubbornness</i>	<i>bullheadedness</i> , <i>inflexibility</i> , <i>obduracy</i> , <i>obstinacy</i> , <i>stubbornness</i>

Table 3: Comparing thesaurus clusters with clusters emerging from the survey study (English). Terms in italic are present in both clusters.

IH-Synonyms German		
	Thesaurus Cluster Members	Survey Cluster Members
Cluster 1: Sensibility	<i>Achtsamkeit, Anstand, Aufmerksamkeit, Diskretion, Empfindsamkeit, Forscherdrang, Interesse, Neugier, Umsichtigkeit, Verständnis, Vorsicht, Wissbegierde, Wissensdurst, Zurückhaltung</i>	<i>Achtsamkeit, Empfindsamkeit, Integrität, Umsichtigkeit, Vorsicht</i>
Cluster 2: Unpretentiousness	<i>Anspruchslosigkeit, Bescheidenheit, Einfachheit, Integrität, Schlichtheit</i>	<i>Anspruchslosigkeit, Anstand, Bescheidenheit, Diskretion, Einfachheit, Schlichtheit, Zurückhaltung</i>
Cluster 3: Knowledge	<i>Intelligenz, Scharfsinn, Urteilskraft, Weisheit</i>	<i>Aufmerksamkeit, Forscherdrang, Intelligenz, Interesse, Neugier, Scharfsinn, Urteilskraft, Verständnis, Weisheit, Wissbegierde, Wissensdurst</i>
IH-Antonyms German		
Cluster 1: Overrating (1)	<i>Egoismus, Egozentrismus, Ichbezogenheit, Ichsucht, Selbstbezogenheit, Selbstverliebtheit</i>	<i>Egoismus, Egozentrismus, Eitelkeit, Ichbezogenheit, Ichsucht, Selbstbezogenheit, Selbstverliebtheit</i>
Cluster 2: Overrating (2)	<i>Anmassung, Arroganz, Aufgeblasenheit, Aufschneiderei, Dreistigkeit, Eingebildetheit, Eitelkeit, Geringschätzung, Grossspurigkeit, Prahlerei, Überheblichkeit, Vermessenheit, Wichtigtuerei</i>	<i>Anmassung, Arroganz, Aufgeblasenheit, Aufschneiderei, Dreistigkeit, Eingebildetheit, Grossspurigkeit, Prahlerei, Überheblichkeit, Vermessenheit, Wichtigtuerei</i>
Cluster 3: Underrating other	<i>Befangenheit, Einseitigkeit, Engstirnigkeit, Schüchternheit, Ungerechtigkeit, Voreingenommenheit, Vorurteil, Zaghaf-tigkeit</i>	<i>Geringschätzung, Ungerechtigkeit, Voreingenommenheit, Vorurteil,</i>
Cluster 3': Underrating self	<i>Befangenheit, Schüchternheit, Zaghaf-tigkeit</i>	<i>Befangenheit, Schüchternheit, Zaghaf-tigkeit</i>
Cluster 4: Dogmatism	<i>Dickköpfigkeit, Halsstarrigkeit, Inflexibilität, Rechthaberei, Sturheit, Verbohrtheit</i>	<i>Dickköpfigkeit, Einseitigkeit, Engstirnigkeit, Halsstarrigkeit, Inflexibilität, Rechthaberei, Sturheit, Verbohrtheit</i>

Table 4: A comparison of thesaurus clusters with clusters emerging from the survey study (German). Terms in italic are present in both clusters.

In order to check whether the survey clusters found using only the frequent patterns are “genuine”, we calculated the cluster densities provided by the participants as follows: Each of the frequent patterns (e.g. the 62 distinct patterns of English synonyms) was attributed to that survey cluster with which it shares the highest number of terms. For example, 4 out of 5 terms of the frequent pattern {circumspection, curiosity, inquiring mind, inquisitiveness, interest, thirst for knowledge} belong to cluster 3 (knowledge), such that the cluster density rating given for this pattern was attributed to cluster 3. Furthermore, we calculated the mean pattern length for all patterns of each cluster, as it is a priori more likely that longer patterns are rated as less dense, since adding more terms increases the likelihood that the pattern is conceived as heterogeneous. When comparing the patterns that belong to the three survey clusters found with all other patterns, we find in several cases (English: 3 cases, German: 5 cases) displaying one of two features. Either the density of the identified clusters is significantly higher (despite the fact that the pattern lengths are comparable or even longer) or the cluster densities are comparable despite much longer mean pattern length in the identified clusters. This is a strong argument that the clusters identified are indeed perceived to be genuine by the participants. This is confirmed in most cases by the names that the participants gave to the clusters: the three most frequent names match very precisely our naming of the dimensions based on the thesaurus analysis (in two English survey clusters the number of entries was too low to see a clear result regarding naming). Recall that the participants were totally free in how to name their groups and the way how to arrange the words. The fact that our a priori naming of clus-

ters in the semantic study overlaps significantly with the naming of clusters by participants in the empirical study suggests that these clusters are generally recognized as such.

IH-Synonyms English				
	Cluster density	Mean pattern length	3 most frequent names	Fraction of most frequent names
Cluster 1: Sensibility	3.88	3.25***	[Number too low]	-
Cluster 2: Unpretentiousness	2.95	7.02	Humility, modesty, personality	47.6%
Cluster 3: Knowledge	2.54***	7.62	Knowledge, intelligence, curiosity	47.7%
All remaining patterns	3.08	7.16	-	-
IH-Antonyms English				
Cluster 1: Overrating (1)	2.23***	11.86***	Egoism/egocentrism, self-centeredness, narcissism	59.1%
Cluster 2: Overrating (2)	2.91	6.36	Arrogance	54.5%
Cluster 3: Underrating other	2.54	5.40	Bias, prejudice, unfairness	59.0%
Cluster 4: Underrating self	4.09	3.00*	[Number too low]	-
Cluster 5: Dogmatism	2.03***	5.69	Stubbornness	62.9%
All remaining patterns	2.80	6.13	-	-

Table 6: Cluster densities and mean pattern length of survey clusters, as well as the three most frequent names (English): By ‘names’ we refer to the *precise* wording used by participants, counting only variants like singular/plural or nouns/adjectives as a single name.

IH-Synonyms German				
	Cluster density	Mean pattern length	3 most frequent names	Fraction of most frequent names
Cluster 1: Sensibility	2.71	6.29	Sozialkompetenz, Empathie, Achtsamkeit	54.2%
Cluster 2: Unpretentiousness	2.49***	5.24***	Bescheidenheit, Einfachheit, Zurückhaltung	38.7%
Cluster 3: Knowledge	2.32***	7.00	Wissen(-sdurst/schaft), Forschen, Neugier	52.3%
All remaining patterns	3.06	7.18	-	-
IH-Antonyms German				
Cluster 1: Overrating (1)	1.88***	6.82***	Egoismus, Egozentrismus, Ichbezogenheit	46.6%
Cluster 2: Overrating (2)	2.38	8.51***	Vorurteil Voreingenommenheit, Ungerechtigkeit	46.6%
Cluster 3: Underrating other	2.51	4.01***	Arroganz, Überheblichkeit, Angeberei	31.0%
Cluster 4: Underrating self	2.46	3.15***	Schüchternheit, Unsicherheit, Zurückhaltung	45.3%
Cluster 5: Dogmatism	2.27***	7.57***	Sturheit, Inflexibilität, Engstirnigkeit	42.7%
All remaining patterns	2.58	6.23	-	-

Table 6: Cluster densities and mean pattern length of survey clusters, as well as the three most frequent names (German): By ‘names’ we refer to the *precise* wording used by participants, counting only variants like singular/plural or nouns/adjectives as a single name.

5. Discussion and Conclusion

From these results, there are three points we wish to draw out for discussion. First, there is the matter of what the clusters represent. In the antonyms map, we take each cluster to represent a distinct vice, i.e., a different way one can fail to be intellectually humble. For the synonyms, however, two possibilities exist. It might be that each cluster represents a distinct trait, all three of which go by the same name of ‘intellectual humility’. Opposing this semantic diversity thesis is the alternate interpretation that sees each cluster representing a different facet of the single trait of intellectual humility.

Second, consider the merging of the terms belonging to the synonym-based unpretentiousness dimension and antonym-based underrating-self dimension in the combined map. We see three possible interpretations. It might be that the discreet aspect of intellectual humility is essentially akin to underrating oneself. Snow (1995) and Taylor (1985) both argue that humility essentially involves recognizing one’s low status or personal faults. Likewise, Julia Driver (2001, p. 21; see also Sidgwick 1907/1962, p. 334) suggests that the humble person “is disposed to underestimate self-worth to some limited extent, in spite of the available evidence.” If this is right, then either the discreet aspect of humility is more of a vice than a virtue, or the underrated aspect of humility’s antonyms is more of a virtue than a vice. Either way, the valence of one or both of these semantic clusters may need to change.

This is an exploratory study, so any conclusions we draw are speculative. That said, another plausible interpretation of the overlap of synonyms and antonyms draws on Nietzsche’s (1887 / 1967) claim that Christianity brought about a “revaluation of values.” In ancient European societies, pride was generally considered a virtue and humility a vice (arrogance was also considered a vice by, for instance, Aristotle). Christianity reversed this valuation, elevating humility and denigrating pride. We moderns have inherited both value-schemas in a kind of palimpsest, in which the old vice is overlaid by the new virtue.

Finally, there might be two different traits picked out by these clusters – one a virtue and the other a vice – that are behaviorally similar enough that they are easily conflated. Someone who underrates herself will behave very similarly to a discreet person. They will both not regularly speak up about controversial topics, in praise of themselves, or for their own rights and entitlements, making it difficult to differentiate them behaviorally. There could, however, be an underlying psychological difference that typically goes unobserved. The unpretentious person may not often attend to evaluating herself, but when she does so, she does it accurately. One who underrates herself, however, may pay significant attention to her own merits, but regularly devalue them. Further research on the behavioral and psychological aspects of intellectual humility and its contraries may help to answer this question.

The final point about the overlap of IH’s synonyms and antonyms relates back to the Big Six personality inventory (Ashton et al. 2004; Saucier 1997). As mentioned, the H factor is meant to represent facets of personality related to honesty and humility. The 100-item revised version measures the participant’s humility (specifically her modesty) by having her indicate (dis)agreement with statements such as “I am an ordinary person who is no better than oth-

ers.” We worry that the Big Six therefore includes in its H dimension items that are better understood as contrary to humility, not allied with or constitutive of it.

Regarding the cross-cultural recognition of intellectual humility and its opposites, our results are salubrious. Naturally, there was not a perfect match between English and German in either the thesaurus study or the empirical study. Nevertheless, similar clusters emerged in both languages. This suggests that, at least in modern European languages and cultures, conceptions of intellectual humility, related virtues, and opposing vices are largely shared. This corresponds to a growing consensus in philosophy and cross-cultural psychology that there is more variance within cultures than between them in basic values.⁴

Finally, we note several important shortcomings of our study. First, the nature of our data does not allow us to determine whether the similarities found reflect merely pragmatic use of words within (American) English and German, or whether they actually represent relations between mental concepts. For example, a hypothesis one could have is that the differences in semantic neighborhoods of terms point to differences in the accessibility of the related mental concepts (Higgins, 1996). This requires additional empirical research that is beyond the scope of this work. Second, the similarity measure used may be misleading in that we did not take into account differences in usage frequency of terms. For example, two pairs of terms may share an overlapping set of the same relative size – but in one case the overlap may consist of terms that are very frequently used as synonyms for the terms under investigation, whereas in the other case, the word bags share synonyms that are only very rarely used. It would be plausible to treat these two cases differently. However, this would again need additional empirical data on how frequent certain terms are used as synonyms. To diminish this problem, services like *Google Ngram* (<https://books.google.com/ngrams>) could be used to obtain proxy data regarding term frequency in synonymy relations. Third, the survey study may be considered insufficiently representative for matching the general language practice regarding similarity assessment. This is more problematic for the German survey that approached Swiss-German speaking people, i.e. a rather small sample within the German-speaking world. Furthermore, *Amazon Mechanical Turk* is known to have a higher demographic diversity compared to university student and staff populations, although some differences to the general population remain (Ross et al., 2010). Fourth, the differences found may also reflect to some degree differences of the thesaurus databases themselves. The German database generated in the mean more synonyms per term compared to the English database. Whether this is indeed a property of the language itself or just of the databases, should be investigated further.

Acknowledgement: We thank Daniel Lapsley, ACE Collegiate Professor and Chair of the Department of Psychology of the University of Notre Dame, and Paul C. Stey, Department of Psychology of the University of Notre Dame, for their input to this research. This work was made possible through a grant from the Thrive Center at Fuller Theological Seminary.

⁴ For philosophical views, see Alfano (2016, chapter 5), Moody-Adams, (1997), and Wong (2006). For psychological views, see Fischer et al. (2010) and Shweder (2012).

References

- Alfano M (2015) Ramsifying virtue theory. In M. Alfano (ed.), *Current Controversies in Virtue Theory*. Routledge, Oxford, pp. 123-135
- Alfano M (2016) *Moral Psychology: An Introduction*. Politi, London
- Allport G, Odbert H (1936) Trait-names: A Psycho-lexical Study. *Psychological Monographs* 47(1), i-171
- Ashton M, Lee K, Perugini M, Szarota P, de Vries R, Di Blas L, Boies K, De Raad B (2004) A six-factor structure of personality-descriptive adjectives: Solutions from psycholexical studies in seven languages. *Journal of Personality and Social Psychology* 86:2,356-366
- Christen M, Alfano M, Bangerter E, Lapsley D (2013) Ethical issues of ‘morality mining’: When the moral identity of individuals becomes a focus of data-mining. In Rahman H, Ramos I (eds.), *Ethical Data Mining Applications for Socio-Economic Development*. Hershey, IGI Global, pp. 1-21
- Christen M, Robinson B, Alfano M (2014) The semantic space of intellectual humility. In Herzig A, Lorini E (eds.), *Proceedings of the European Conference on Social Intelligence*. Pp. 40-49
- Driver J (2001) *Uneasy Virtue*. New York, Cambridge University Press
- Fairbanks SJ (2010) Environmental goodness and the challenge of American culture. *Ethics and the Environment* 15(2):79-102
- Fischer R, Vauclair C-M, Fontaine J, Schwartz S (2010) Are individual-level and country-level value structures different? Testing Hofstede’s legacy with the Schwartz value survey. *Journal of Cross-Cultural Psychology* 41:135-51
- Fricker M (2007) *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford, Oxford University Press
- Galton F (1884) *Measurement of Character*. *Fortnightly Review* 36: 179-185
- Harris R (1973) *Synonymy and Linguistic Analysis*. Toronto, University of Toronto Press
- Higgins ET (1996). Knowledge activation: Accessibility, applicability and salience. In Higgins ET, Kruglanski AW (eds.) *Social psychology: Handbook of basic principles*. New York, Guilford, pp. 133-168
- Hill T (1999) Ideals of human excellence and preserving the natural environment. In Gruen L, Jamieson D (eds.), *Reflecting on Nature: Readings in Environmental Philosophy*. New York

- Hüllen W (2004) *A History of Roget's Thesaurus*. Oxford, Oxford University Press
- Jamieson D (2007) When utilitarians should be virtue theorists. *Utilitas* 19(2):160-182
- Jones SK (1986) *Synonymy and Semantic Classification*. Edinburgh, Edinburgh University Press
- Lewis D (1966) An argument for the identity theory. *The Journal of Philosophy* 63(1):17-25
- Lewis D (1970) How to define theoretical terms. *The Journal of Philosophy* 67(13):427-46
- Lewis D (1972) Psychophysical and theoretical identifications. *Australasian Journal of Philosophy* 50:249-258
- Moody-Adams M (1997) *Fieldwork in Familiar Places: Morality, Culture, and Philosophy*. Cambridge, Harvard University Press
- Nietzsche F (1887 / 1967) *The Genealogy of Morals*. Trans. W. Kaufmann. New York, Random House
- Ostendorf F (1990) *Sprache und Persönlichkeitsstruktur: Zur Validität des Fünf-Factoren-Modells der Persönlichkeit*. Regensburg, S. Roderer Verlag
- Ott T, Eggel T, Christen M (2014) Generating Low-Dimensional Denoised Embeddings of Nonlinear Data with Superparamagnetic Agents. *Proceedings of the 2014 International Symposium on Nonlinear Theory and its Applications (NOLTA)*, Lucerne, Switzerland, September 14-18
- Ott T, Kern A, Steeb W-H, Stoop R (2005) Sequential Clustering: Tracking Down the Most Natural Clusters. *Journal of Statistical Mechanics: theory and experiment*: P11014
- Peabody D, Goldberg L (1989) Some determinants of factor structures from personality-trait descriptors. *Journal of Personality and Social Psychology* 57(3):552-567
- Quine WVO (1951) Two dogmas of empiricism. *The Philosophical Review* 60:20-43
- Quine WVO (1960) *Word and Object*. Cambridge MA, MIT Press
- Ramsey F (1931) Theories. In Braithwaite RB (ed.), *The Foundations of Mathematics*. London, Routledge & Kegan Paul
- Roberts R, Wood J (2007) *Intellectual Virtues: An Essay in Regulative Epistemology*. Oxford, Oxford University Press

Ross J, Irani I, Silberman M, Zaldivar A, Tomlinson B (2010) Who are the Crowdworkers? Shifting Demographics in Amazon Mechanical Turk. In Proceedings of the Conference on Human Factors in Computing Systems (Atlanta, GA, April 10-15), pp. 2863-2872

Saucier G (1997) Effects of variable selection on the factor structure of person descriptors. *Journal of Personality and Social Psychology* 73(6):1298-1312

Shweder R (2012) Relativism and universalism. In Fassin D (ed.), *A Companion to Moral Anthropology*. Oxford, Wiley, pp. 85-102

Sidgwick H (1907/1962) *The Methods of Ethics*, 7th ed. Chicago, University of Chicago Press

Spiegel JG (2012) Open-mindedness and intellectual humility. *Theory and Research in Education* 10:27-38

Snow N (1995) Humility. *The Journal of Value Inquiry* 29(2):203-216

Taylor G (1985) *Pride, Shame and Guilt*. Oxford, Clarendon Press

Trier J (1931) *Der deutsche Wortschatz im Sinnbezirk des Verstandes: Die Geschichte eines sprachlichen Feldes*. Vol. 1: Von den Anfängen bis zum Beginn des 13. Jahrhunderts. Heidelberg, Quelle und Meyer

Wiggins JS (1973) *Personality and Prediction: Principles of Personality Assessment*. Reading, MA, Addison-Wesley

Wong D (2006) *Natural Moralities: A Defense of Pluralistic Relativism*. Oxford, Oxford University Press