

## Research article

## Moral universals: A machine-reading analysis of 256 societies

Mark Alfano<sup>a,\*</sup>, Marc Cheong<sup>b</sup>, Oliver Scott Curry<sup>c</sup><sup>a</sup> Macquarie University, Australia<sup>b</sup> University of Melbourne, Australia<sup>c</sup> University of Oxford, UK

## ARTICLE INFO

## Keywords:

Morality  
Cooperation  
Ethnography  
Universals  
Natural language processing  
LIWC

## ABSTRACT

What is the cross-cultural prevalence of the seven moral values posited by the theory of “morality-as-cooperation”? Previous research, using laborious hand-coding of ethnographic accounts of ethics from 60 societies, found examples of most of the seven morals in most societies, and observed these morals with equal frequency across cultural regions. Here we replicate and extend this analysis by developing a new Morality-as-Cooperation Dictionary (MAC-D) and using Linguistic Inquiry and Word Count (LIWC) to machine-code ethnographic accounts of morality from an additional 196 societies (the entire Human Relations Area Files, or HRAF, corpus). Again, we find evidence of most of the seven morals in most societies, across all cultural regions. The new method allows us to detect minor variations in morals across region and subsistence strategy. And we successfully validate the new machine-coding against the previous hand-coding. In light of these findings, MAC-D emerges as a theoretically-motivated, comprehensive, and validated tool for machine-reading moral corpora. We conclude by discussing the limitations of the current study, as well as prospects for future research.

## 1. Introduction

According to the theory of “Morality-as-Cooperation”, morality is a collection of cooperative rules that help humans work together, keep the peace, and promote the common good [1]. There are many types of cooperation, hence many types of morality, including: family values, group loyalty, reciprocity, heroism, deference, fairness, and property rights. Previous research, using laborious hand-coding of ethnographic accounts of ethics from 60 societies, found examples of most of these seven moral values in most societies, and observed these morals with equal frequency across cultural regions [2]. Here we extend this analysis by developing the Morality-as-Cooperation Dictionary (MAC-D) and using LIWC to machine-code ethnographic accounts of ethics from an additional 196 societies.

The paper proceeds as follows. First, we outline the theory of “Morality-as-Cooperation”, and summarise previous work testing its predictions regarding the cross-cultural prevalence of distinct moral values. Second, we introduce the natural language processing Linguistic Inquiry and Word Count (LIWC) approach, and we describe our methodology for developing MAC-D. Third, we use MAC-D to estimate the cross-cultural prevalence of MAC’s seven moral values in a total of 256 societies; and we compare the results of the hand- and machine-coding of the original 60 societies. Finally, we discuss the strengths and limitations of the current research, and consider the prospects for further research using machines to code moral corpora.

\* Corresponding author. 25B Wally’s Walk, NSW, 2109, Australia.  
E-mail address: [mark.alfano@gmail.com](mailto:mark.alfano@gmail.com) (M. Alfano).

<https://doi.org/10.1016/j.heliyon.2024.e25940>

Received 17 March 2023; Received in revised form 30 January 2024; Accepted 5 February 2024

Available online 10 February 2024

2405-8440/© 2024 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

## 2. The theory of morality-as-cooperation

Recent research suggests that the function of morality is to promote cooperation [1,3 p. 40, 4 p. 800, 5 p. 59, 6 p. 1, 7 p. 231]. According to this theory of “Morality-as-Cooperation” (MAC), humans face, and have faced, a range of different nonzero-sum problems of cooperation, and have evolved and invented a range of solutions to them. These cooperative solutions take a variety of forms, including character traits, strategies, dispositions, behaviours, rules, norms, institutions, and technologies. Together, they motivate cooperative behaviour and provide the criteria by which we judge the behaviour, attitudes, and traits of ourselves and others.<sup>1</sup> And it is this collection of cooperative solutions that philosophers and others have called morality [1].

What’s more, because there are many types of cooperation, MAC predicts that there will be many types of moral values. For example, drawing on evolutionary game theory, Curry [1] has argued that there are (at least) seven distinct types of cooperation: (1) the allocation of resources to kin; (2) coordination to mutual advantage; (3) social exchange; and conflict resolution through contests featuring (4) hawkish displays of dominance and (5) dovish displays of deference; (6) division of disputed resources; and (7) recognition of prior possession. And each of these types of cooperation gives rise to a corresponding type of morality: (1) family values, (2) group loyalty, (3) reciprocity, (4) heroism, (5) deference, (6) fairness, and (7) property rights.

### 1. Family Values

Kin selection explains why we love and care for our families, why parents feel a duty of care to their offspring, why we feel a special obligation to help our extended families, and why we abhor incest [3,4].

### 2. Group Loyalty

Mutualism explains why we coordinate our activities to pursue projects of mutual interest, why we form groups, clubs, and coalitions (there is strength and safety in numbers), why we value these groups, their members, and our membership in them, why we adopt local norms and conventions, why we feel a special obligation to come to the aid of fellow group members, and why we value loyalty, unity, and solidarity [5,6].

### 3. Reciprocity

Social exchange explains why we seek opportunities for mutually-beneficial trade, and why we feel we ought to trust others, return favours (positive reciprocity), keep promises, pay debts, fulfill contracts, be grateful for favours received, feel guilty for favours not returned, avenge injuries (negative reciprocity), punish cheaters, apologise for causing injuries, and forgive those who apologise [7,8].

### 4. & 5. Heroism and Deference

Conflict resolution explains why we minimise the mutual costs of disputes by engaging in ritual contests: why we proudly display cues of power and high status (‘excellences’, including bravery and generosity); and, when bested, why we express humility, and respect, defer, and submit to our superiors [9,10].

### 6. Fairness

Conflict resolution also explains why we resolve disputes over divisible resources by dividing or sharing them (often resulting in equal shares), and hence why we feel an obligation to negotiate, seek a compromise, and be fair [11,12].

### 7. Property rights

Conflict resolution also explains why we resolve disputes by recognising prior possession, hence why we feel we ought to respect others’ property and territory, and refrain from theft [13].

MAC’s theory-derived taxonomy of moral domains extends beyond, and provides a principled and more comprehensive alternative to, those provided by earlier more ad hoc taxonomies, such as Moral Foundations Theory (MFT; [14]). For example, MAC includes domains dedicated to family values, reciprocity, heroism and property rights, which are missing from MFT. As such, MAC is able to predict and explain a broader array of moral phenomena [1,2,15,16].

<sup>1</sup> MAC is primarily a theory of the function of morality, specifying that whatever tends to contribute to helping humans cooperate in non-zero-sum interactions that are common in social life will be considered morally good, and that whatever tends to undermine such cooperative interactions will be considered morally bad. For this reason, MAC is indifferent to what structures realize this function; presumably, many different things do so in different contexts, from innate sentiments, to inherited social and cultural norms, to constructed institutions and technologies. For this reason, MAC is consistent with many different ways of enacting and embodying moral values. Just as there are many different kinds of eating utensils that share the same functional role, so (according to MAC) there are many local implementations of cooperative strategies that share the same functional role. For additional philosophical background on this point, see [8–10, 11 ch. 5, 12].

MAC makes a number of predictions that have been supported by previous research. For example, MAC predicts that: there will be as many different types of morality as there are types of cooperation (and their combinations) [15]; that each of these distinct types of morality will be the product of genetically-distinct, domain-specific psychological mechanisms [17]; that each type of morality will emerge spontaneously during development [18]; that each of these types of cooperation will be considered morally relevant and will constitute a distinct facet or factor [16]; and, most relevant to the current research, that each of these types of cooperation will be considered morally good, in all societies [2].

In order to test this cross-cultural prediction, Curry and colleagues conducted a content analysis of 600 ethnographic accounts of ethics from 60 societies, using the holocultural method [2,19,20]. The 60 societies represent the 'Probability Sample Files' (PSF), a stratified random sample of well-attested human societies, drawn from the electronic Human Relations Area Files (eHRAF) – an archive of thousands of original, full-text ethnographies. They extracted paragraphs that had been labeled as Ethics or Norms, or that contained one of a set of moral keywords — 3460 paragraphs in total. They coded whether each paragraph contained a morally positive or morally negative reference to the seven types of cooperation (note that they did not use a dictionary for this, but rather a holistic assessment of each paragraph by trained coders). 761 paragraphs contained information about the moral valence of one or more of the seven cooperative behaviours, giving rise to 962 observations of moral valence in total. They reported three main results. First, in 961 out of 962 observations (99.9%), cooperative behaviour had the predicted positive moral valence. (The one exception was a case in which 'open theft' was considered admirable because it displayed dominance.) Second, the seven positively-morally valenced cooperative behaviours were observed in 78%, 73%, 72%, 53%, 58%, 15%, and 90% of societies respectively. And third, all seven positively-morally valenced cooperative behaviours were observed with equal frequency across six cultural regions. Thus, MAC's predictions were supported; these seven types of morality appear to be cross-culturally universal.

One limitation of this study is that, because hand-coding is very labour-intensive, it was impractical to go beyond the 60 PSF societies and investigate the prevalence of morals in the other societies documented by HRAF. Another limitation is that hand-coding for the *presence or absence* of a moral is a coarse-grained measure, which makes it difficult to investigate variation in the degree of endorsement of a given moral in a given society. In order to overcome these limitations, here we develop and deploy a semi-automated approach that harnesses natural language processing tools to detect continuous variation in moral values across societies in the entire ethnographic corpus.

### 3. Developing the MAC-D

In this section, we explain the methodology used to develop the MAC-D sub-dictionaries. We begin by explaining the expert-led initial step, then follow up with the semi-automated and expert-curated step employing WordNet.

#### 3.1. Linguistic Inquiry and Word Count (LIWC) and the LIWClike R package

Recent work has shown it is possible to extract psychological signals from natural language using various digital tools. One prominent approach, Linguistic Inquiry and Word Count (LIWC) — developed by Pennebaker et al. [21] (see also [22]) — works by counting the number of words belonging to various categories in a text. For instance, LIWC combs through a text to examine how many first-person singular pronouns (e.g., *I, me, my*), how many first-person plural pronouns (e.g., *we, us, our*), and how many articles (e.g., *a, an, the*) are used per 100 words. Beyond these so-called *function words*, LIWC has dictionaries for words that indicate complexity of thinking (e.g., *nevertheless, whereas, but*), asking questions (e.g., *who, what, why*), and a range of psychological processes that includes positive and negative emotions, discrete emotions such as anger and sadness, and drives for affiliation, achievement, power, reward, and risk. Over the last decades, Pennebaker and his collaborators have shown that this method can reveal quite a bit about individuals, their relationships, and the groups to which they belong. One useful function of LIWC is the ability to create and share custom dictionaries for categories of interest. The dictionary approach is favored as it “has the advantage of being ‘objective, automated, and transparent’ ...” [23] (as cited in Ref. [24]); this transparency-by-design allows us to “provide a clear explanation as to how and why” [24] certain signals are detected in a particular piece of text. The portability of the LIWC (.dic) file format, and the existence of multiple free and open-source implementations of the LIWC algorithm, are added advantages.

Related approaches have been applied to moral corpora. For example, Alfano, Higgins, & Levernier [25] have advocated taking a lexical approach to guide philosophical investigations into moral values and virtues. In a case study, they extract obituary documents' agent-level descriptions of the deceased that refer to their character traits and relationships, thereby producing a collocation network of virtues, values, and constituents of wellbeing.

Most previous research using LIWC to analyse moral corpora has employed the Moral Foundations Dictionary (MFD) and concentrated on contemporary (mostly American) corpora [14]. For example, Hoover et al. [26] annotated a large corpus of tweets for five moral domains, as did Trager et al. [27] for a Reddit corpus (non-peer-reviewed pre-print). Padfield & Buchanan [28] use MFD to characterise the moral lean of various US media organisations. And Buttrick et al. [29] track the evolution of Anglophone moral language over the course of more than two centuries. However, this MFD-based approach is limited in at least two ways. First, there are underlying problems with Moral Foundations Theory, on which it is based. Moral Foundations Theory, and thus MFD, omits four key moral domains: it has no foundations dedicated to family values, reciprocity, heroism, or property rights. Second, when it comes to the

updated MFD 2.0, the estimates of reliability and validity (see <https://osf.io/xakyw>) have not been peer-reviewed, and the methodology is sparsely documented. Yet a third Moral Foundations dictionary — the extended MFD or eMFD — was recently published [30]. It relies on newer methods but is not directly comparable to the list-based dictionary approach taken here.<sup>2</sup> Future work using these newer methods to develop an extended MAC dictionary or eMAC-D would enable more of an apples-to-apples comparison between dictionaries associated with Moral Foundations Theory and Morality-as-Cooperation.

To address these limitations, we developed the Morality-as-Cooperation Dictionary (MAC-D) for use in either the proprietary LIWC platform, or in open-source packages for both the R statistical computing environment and the Python programming language.<sup>3</sup> Note that, while MAC and MFT are theoretical competitors, both are premised on the idea that morality is all (in the case of MAC) or mostly (in the case of MFT) about cooperation. In order for their predictions to eventually be compared in corpora in future research, MAC-D needs to be established as a potential competitor to MFD and MFD 2.0. We do not make that comparison in the current paper because there is as yet no ground truth coding of Moral Foundations in HRAF or the 60 societies documented in the PSF.

### 3.2. Initial expert-led enumeration of words, stems, and n-grams

Starting in 2020 (and thus before the efflorescence of large language models such as ChatGPT), the authors began constructing MAC-D by brainstorming words, stems, and n-grams that describe when someone was acting in accordance with, or in violation of, the seven basic moral principles constitutive of MAC – love your family, be loyal to your group, return favours, be brave and generous,<sup>4</sup> defer to superiors, divide disputed resources fairly, and respect property. This led to the creation of seven virtue proto-dictionaries and seven vice proto-dictionaries.<sup>5</sup>

Next, we solicited additional words, stems, and n-grams from colleagues (all Ph.D.s in the humanities or social sciences) who were already familiar with MAC, five of whom responded. We took the union of their lists and added those to the fourteen proto-dictionaries we had already enumerated. In sum, for the virtue dictionaries, three to four dozen items were included in these proto-dictionaries, while for each of the vice dictionaries approximately two dozen items were included.

Next, the authors reached consensus on whether the words, stems, and n-grams in the proto-dictionaries had face-valid fit for the categories of MAC as described above. For example, the term ‘own’ could not be included in the property dictionary because, while it can refer to ownership of property, it can also be used in phrases like ‘her own family’. The authors also continued to brainstorm new words, stems, and n-grams that might be added to each proto-dictionary. Some terms were associated with multiple proto-dictionaries (e.g., ‘admire’, which was associated with both heroism and deference). Due to the way LIWC parses tokens, the authors introduced punctuated versions of words and abbreviations (e.g., ‘iou’ and ‘give and take’, versus ‘i.o.u.’ and ‘give-and-take’), as LIWC regards the latter terms as distinct from the former. All told, the virtue proto-dictionaries comprised 371 items (a mix of words, stems, and n-grams), while the vice proto-dictionaries comprised 148 items. These proto-dictionaries represented the expert input of all co-authors and five colleagues.

### 3.3. Expansion via WordNet, followed by culling of false positives

In order to ensure that the proto-dictionaries developed by experts were comprehensive and not too idiosyncratic, we next turned to WordNet.<sup>6</sup> Wordnet is a large lexical database of the English language in which words are grouped into sets of cognitive synonyms (synsets). Each synset expresses a distinct concept and is related to other synsets by various conceptual-semantic and lexical relations. Using WordNet, we aimed to discover additional words that should be added to each of the proto-dictionaries. WordNet has been used as a standard tool for discovering relations between words in natural language processing studies, and has been cited in over 14,000 academic studies.<sup>7</sup>

Given an input word or phrase, WordNet generates synsets that can then be used to query related terms in the following categories: synonyms, antonyms, hyponyms (specific instances for a concept, e.g. ‘pasta’ → ‘fettuccine’), hypernyms (generic categorisations of a concept, e.g. ‘pasta’ → ‘food’), meronyms (sub-components found in a concept, e.g. ‘room’ → ‘ceiling’), holonyms (the larger component of which a concept belongs to, e.g. ‘room’ → ‘building’), and entailments (logical consequences of a word, e.g. ‘eat’ → ‘swallow’).

As WordNet does not work with stems (e.g. ‘commun\*’), in order to generate synsets as above, all word stems were “expanded” by

<sup>2</sup> For recent review of state-of-the-art methods, including some work on morality in corpora, see Ref. [50]; the chapters by Atari & Dehghani and Leone & Niemi are especially relevant, though they do not address the measurement of Morality-as-Cooperation.

<sup>3</sup> LIWC is implemented as a function in the *quanteda* package [51]. Various open-source LIWC implementations exist for Python as well, such as <https://github.com/chbrown/liwc-python>.

<sup>4</sup> Note that Curry et al. [2] coded for heroism = bravery, whereas here we are coding heroism = bravery or generosity, so the results are not directly comparable.

<sup>5</sup> We note the year in which the methodology was developed and deployed because natural language processing and generation have witnessed astonishing advances in the meantime, especially with the public release of ChatGPT (GPT-3.5) at the end of 2022, which we will discuss in Section 5. The current work, however, should not be discounted simply because it took time to complete in the middle of a pandemic and was thus not able to benefit from the guidance of more recent publications, such as [52–54]. We do note, however, that Nicolas and colleagues [52] recommend WordNet expansion and other methods that we independently alighted on and employed in this study.

<sup>6</sup> See <https://wordnet.princeton.edu/> for documentation.

<sup>7</sup> See Ref. [55] for a technical exposition. Citations per Google Scholar as of 31 July 2020.

consulting the NLTK words corpus<sup>8</sup> for terms containing the stem in question (e.g. ‘commun\*’ → ‘communa’, ‘communal’, ‘communism’, etc.). Finally, we queried WordNet for related terms (from all categories above) given each proto-dictionary’s contents, and consolidated them by removing duplicate entries.

This process greatly increased the number of words and n-grams in the proto-dictionaries. The largest virtue proto-dictionary (heroism) contained 776 items, while the smallest virtue proto-dictionary (property) contained 190 items. The largest vice proto-dictionary (group) contained 357 items, while the smallest vice proto-dictionary (family) contained 29 items. However, many of the new items appeared to be false positives. For example, WordNet took the term ‘member’ to refer not only to a member of a group but also to male genitalia, then suggested many synonyms. To handle this problem, two of the co-authors independently coded each term in each proto-dictionary for relevance (1 = is associated with the virtue or vice in question, 0 = is not associated with the virtue or vice in question). Inter-rater reliability was calculated as Cohen’s kappa in the irr R package [31]. For the virtue dictionaries, inter-rater reliability was substantial ( $N = 3038$ ,  $\kappa = 0.70$ ,  $z = 38.8$ ,  $p < .0001$ ). For the vice dictionaries, inter-rater reliability was lower but also substantial ( $N = 1307$ ,  $\kappa = 0.65$ ,  $z = 23.5$ ,  $p < .0001$ ). Because WordNet includes the antonym relation among synsets, we then inspected the words that were included as antonyms in each proto-dictionary. By consensus, we agreed that the antonyms for all virtue and vice proto-dictionaries should be swapped, with the exception of the family proto-dictionaries. This was because the virtue-reciprocity proto-dictionary contained antonyms such as ‘distrust’, ‘impenitent’, and ‘unapologetic’, whereas the vice-reciprocity proto-dictionary contained antonyms such as ‘gratitude’, ‘penitent’, and ‘repentant’. The vice-family proto-dictionary was empty, and the virtue-family proto-dictionary contained words like ‘brother’ (antonym of ‘sister’). The only exception was ‘uncousinly’, which we moved from the virtue-family proto-dictionary to the vice-family proto-dictionary.

During the coding process, all authors also identified a handful of words that wound up in arguably the wrong proto-dictionary; for example ‘reverence’ was located in the virtue-heroism proto-dictionary, but actually fits better in the virtue-deference proto-dictionary. These words were moved into the correct proto-dictionary by consensus.

Finally, all stems left over from the original expert-derived proto-dictionaries were expanded out to all of their potential completions using the method described above. For n-grams, only the last word was expanded. The second co-author noted that certain plurals of less-frequently used terms, such as ‘bloodlines’, do not exist in the NLTK words corpus. To address this problem, we supplemented the word list used for expansion with the Spell Checking Oriented Word Lists (SCOWL) collection, which contains over 670,000 words from multiple varieties of English (UK, US, Canadian, and Australian).<sup>9</sup> Post-expansion, this also resulted in a large number of false positives. For example, from the group dictionary ‘strike break\*’ was expanded to ‘strike breakfast’. These errors were trivially easy to find, so one co-author deleted them. The finalized dictionaries were then converted to LIWC’s native .dic format (see examples in Table 1).

In the next section, we use MAC-D to estimate the prevalence of seven morals in the full ethnographic record of eHRAF.

## 4. MAC-D meets HRAF

### 4.1. Cross-cultural prevalence

We began by extracting paragraphs labeled *Ethics* from eHRAF’s ethnographic archive covering 331 societies (see Fig. 1).<sup>10</sup> This resulted in a total of 1,620,644 words, from 9653 paragraphs, from 1389 documents, covering 256 societies, from 8 cultural regions, following 9 subsistence strategies.<sup>12</sup> (75 societies had no paragraphs labeled *Ethics*).

We then used MAC-D to calculate the percentage of words in each paragraph that was present in each of the seven sub-dictionaries. (A score of ‘10%’ for ‘heroism’, for example, means that one in ten of the words in a given paragraph are in the heroism dictionary.)<sup>13</sup>

The means and standard deviations for each moral domain are shown in Table 2. Because the distribution was heavily skewed (the median and modal percentage for all morals was 0%), we proceeded with non-parametric statistics.

Next we investigated whether these percentages varied between the previously coded PSF60 societies, and the additional 256 non-PSF60 societies. Mann-Whitney U tests revealed that the percentages between these societies were sometimes significantly, but never substantially, different (effect sizes were small,  $r_{bs} \leq 0.08$ ; Table 2).

<sup>8</sup> NLTK is a popular Natural Language Processing toolkit for Python [56]. A specific list of common English words is available in NLTK as the eponymous words corpus. The technical documentation is available at <https://www.nltk.org/>.

<sup>9</sup> SCOWL is an extensive collection of wordlists for the GNU *Aspell* open-source spell checker, which has a simple wordlist generation tool at <http://app.aspell.net/create>. *Aspell* is used in many open-source and commercial software packages. See Ref. [57].

<sup>10</sup> Note that Curry et al. [2] used paragraphs labeled *Ethics*, *Norms* and those that were picked up using specific search terms. Here, to streamline the analysis we focus only on paragraphs labeled *Ethics*.

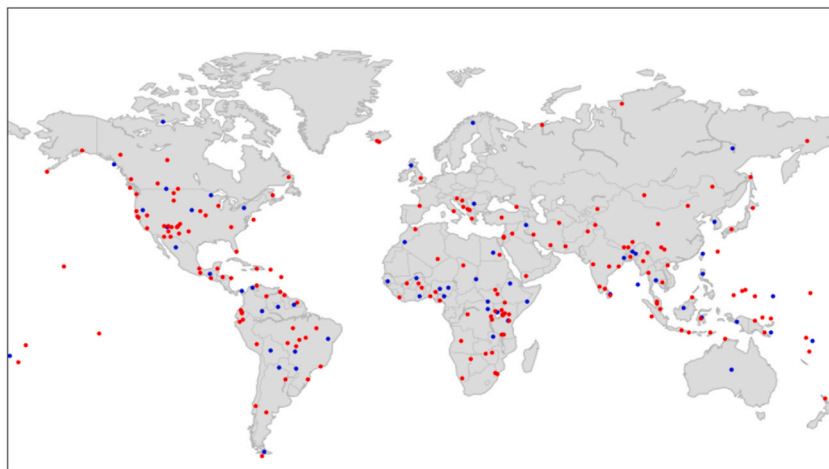
<sup>11</sup> For completeness, a list of data sources used in Basemap cartography is available at <https://matplotlib.org/basemap/users/geography.html>.

<sup>12</sup> The number of paragraphs from each cultural region were as follows: Africa ( $n = 2006$ ), Asia ( $n = 1705$ ), Europe ( $n = 868$ ), Middle America and the Caribbean ( $n = 371$ ), Middle East ( $n = 580$ ), North America ( $n = 2809$ ), Oceania ( $n = 764$ ), South America ( $n = 550$ ). The number of paragraphs from societies with each subsistence strategy were as follows: agro-pastoralists ( $n = 1603$ ), commercial economy ( $n = 816$ ), horticulturalists ( $n = 1017$ ), hunter-gatherers ( $n = 932$ ), intensive agriculturalists ( $n = 3615$ ), not assigned ( $n = 1$ ), other subsistence combinations ( $n = 967$ ), pastoralists ( $n = 455$ ), primarily hunter-gathers ( $n = 247$ ).

<sup>13</sup> Note that because the different sub-dictionaries have different numbers of words with different base rates, we cannot straightforwardly compare the relative frequency of different morals in the same paragraph (or society or region), but we can compare the relative frequency of the same moral between different paragraphs (or societies or regions).

**Table 1**  
Examples from the 14 MAC-D sub-dictionaries.

dictionary	<i>n</i>	examples
<i>virtue</i>		
family	276	clan, cousinly, nanna, paternoster, supermom
group	477	alliance, cadres, collectivised, guilds, shared fate
reciprocity	331	a favor, give-and-take, iou, kept her promise, reciprocated
heroism	454	admirably, brassy, gallant, medaled, venerably
deference	258	abasement, fealty, obeisant, respects, second fiddle
fairness	123	bargaining, equal share, equalitarian, mediator, split the difference
property	118	birthright, first come first served, ownership, prior claim, stake your claim
<i>vice</i>		
family	53	absent father, child abandonment, deadbeat dad, helicopter parent, incest
group	259	apostasy, bilks, deviant, heretical, renegade
reciprocity	109	break her promise, free riders, renege, unforgivably, untrustworthy
heroism	195	cheapskate, cowardly, dishonor, miser, shame
deference	126	above his station, deserter, insubordinately, renegade, sore loser
fairness	35	cronyism, favoritism, inequality, play favorites, unfair
property	239	bandit, biopiracy, carjack, defraud, speculation



**Fig. 1.** Global distribution of ethnographies from eHRAF, across 238 out of the 256 cultures. Blue dots indicate PSF 60 societies while red dots indicate the rest of the cultures. Note that 18 cultures do not have location metadata, due to e.g., spanning multiple field sites across a continent. Location metadata courtesy of The Human Relations Area Files, Inc. (Map generated using the Basemap<sup>11</sup> package for Python).

**Table 2**  
Moral words (%) by MAC domain.

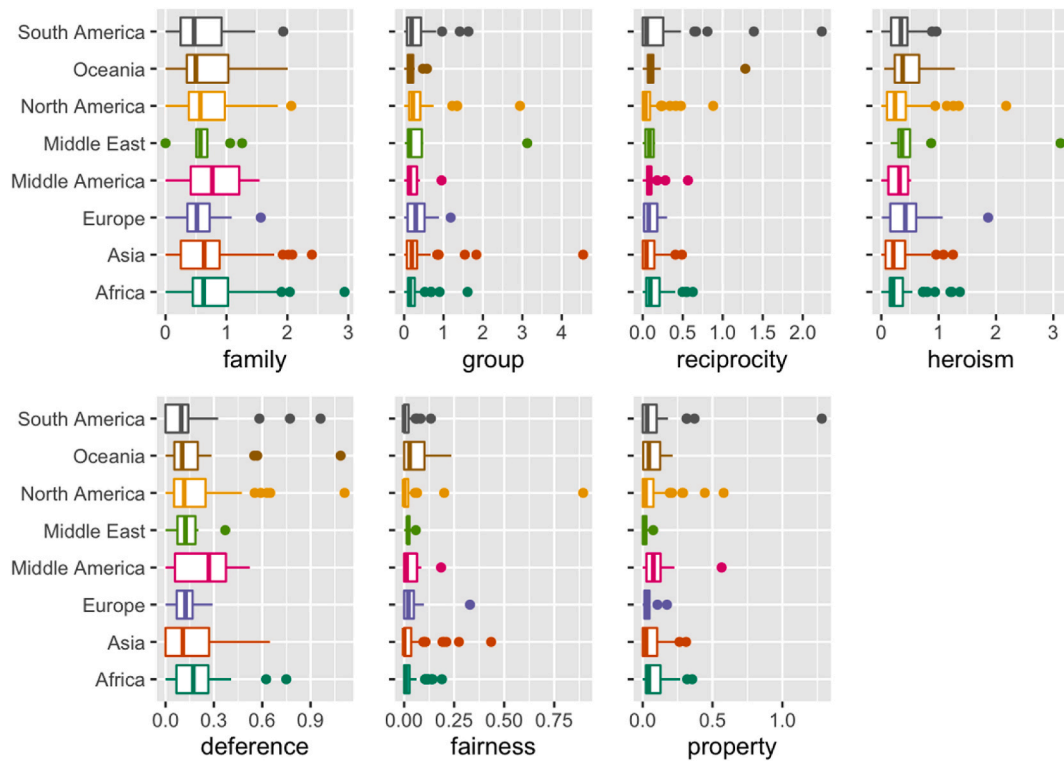
		Family	Group	Reciprocity	Heroism	Deference	Fairness	Property
<b>Total (n=9653)</b>	<b>Mean</b>	0.92	0.27	0.14	0.43	0.22	0.04	0.15
	<b>SD</b>	1.77	0.70	0.46	0.94	0.65	0.22	0.64
<b>PSF (n=2436)</b>	<b>Mean</b>	0.80	0.23	0.14	0.52	0.27	0.04	0.13
	<b>SD</b>	1.48	0.57	0.44	0.97	0.76	0.23	0.54
<b>Non-PSF (n=7217)</b>	<b>Mean</b>	0.97	0.29	0.14	0.40	0.20	0.04	0.15
	<b>SD</b>	1.86	0.74	0.47	0.93	0.61	0.21	0.67
	<b>U</b>	8490000	8530000	8780000	8050000	8480000	8770000	8780000
	<b>p</b>	.01	.01	.87	<.001	<.001	.71	.90
	<b>r<sub>bs</sub></b>	0.034	0.029	0.001	0.084	0.035	0.002	0.001

We then investigated whether these percentages varied across cultural regions. Kruskal-Wallis tests revealed that the percentages of moral words were significantly, but not substantially, different across cultural regions (effect sizes were small:  $0.004 \leq \epsilon^2 \leq 0.056$ ; Table 3; Fig. 2).

Finally we investigated whether these percentages varied across subsistence strategy. Kruskal-Wallis tests revealed that the percentages of moral words were significantly, but not substantially, different across strategies (effect sizes were small:  $0.0007 \leq \epsilon^2 \leq 0.02$ ; Table 4; Fig. 3).

**Table 3**  
Moral Words (mean %) x MAC Moral Domain x Region.

Region	Family	Group	Reciprocity	Heroism	Deference	Fairness	Property
Asia	1.13	0.25	0.16	0.34	0.25	0.06	0.13
Europe	0.83	0.57	0.21	0.91	0.18	0.05	0.12
Africa	1.04	0.23	0.14	0.44	0.26	0.04	0.15
Middle East	0.77	0.20	0.15	0.45	0.15	0.03	0.08
North America	0.88	0.27	0.09	0.27	0.14	0.02	0.19
Middle America and the Caribbean	0.92	0.22	0.11	0.43	0.34	0.04	0.13
Oceania	0.71	0.19	0.14	0.61	0.38	0.04	0.13
South America	0.67	0.30	0.18	0.54	0.13	0.02	0.15
$\chi^2$	132.30	283.60	135.30	535.30	163.40	82.50	36.80
df	7	7	7	7	7	7	7
p	<.001	<.001	<.001	<.001	<.001	<.001	<.001
$\epsilon^2$	0.014	0.029	0.014	0.055	0.017	0.009	0.004



**Fig. 2.** Distribution of moral words across cultural regions.

**Table 4**  
Moral Words (mean %) x MAC Moral Domain x Subsistence Strategy.

Subsistence Strategy	Family	Group	Reciprocity	Heroism	Deference	Fairness	Property
intensive agriculturalists	1.26	0.32	0.14	0.38	0.21	0.04	0.16
hunter-gatherers	0.68	0.31	0.13	0.50	0.18	0.05	0.15
pastoralists	0.61	0.19	0.10	0.78	0.28	0.05	0.16
agro-pastoralists	0.64	0.23	0.17	0.34	0.14	0.02	0.16
horticulturalists	0.74	0.22	0.13	0.38	0.21	0.03	0.15
other subsistence combinations	0.63	0.19	0.15	0.61	0.30	0.04	0.15
commercial economy	1.13	0.31	0.12	0.44	0.34	0.04	0.04
not assigned	0.00	3.13	0.00	3.13	0.00	0.00	0.00
primarily hunter-gatherers	0.54	0.30	0.13	0.51	0.13	0.03	0.19
$\chi^2$	214.52	88.17	7.63	167.69	107.01	27.63	71.18
df	8	8	8	8	8	8	8
p	<.001	<.001	0.47	<.001	<.001	<.001	<.001
$\epsilon^2$	0.022	0.009	0.001	0.017	0.011	0.003	0.007

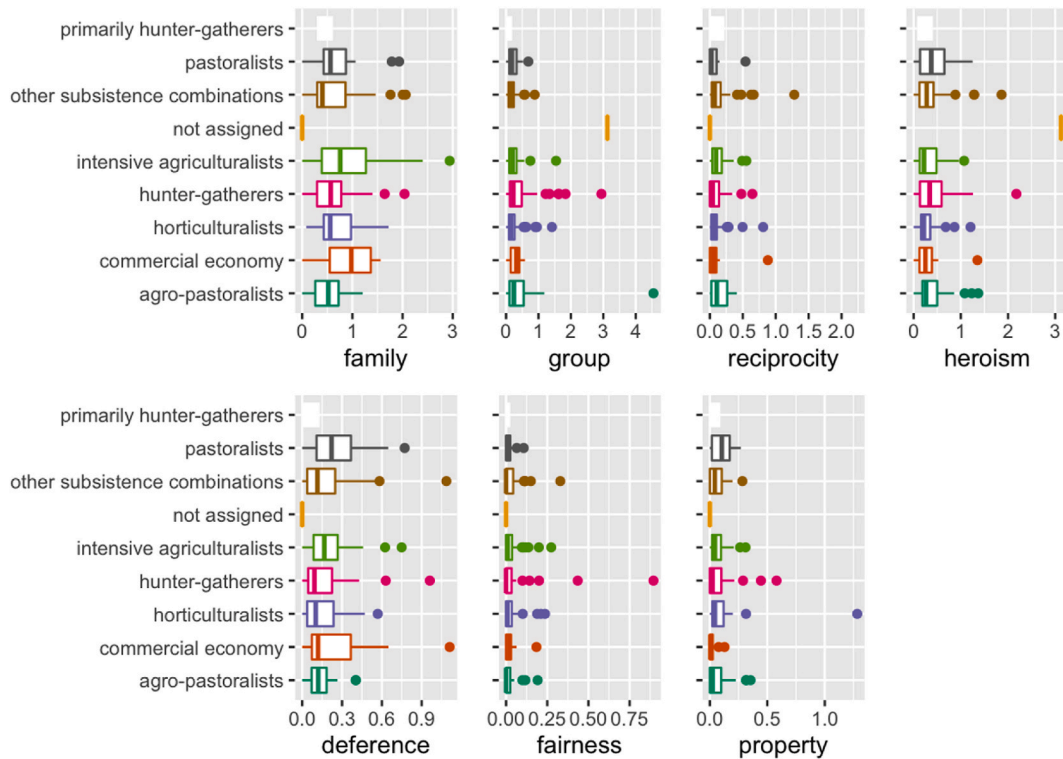


Fig. 3. Distribution of moral words across subsistence strategy.

4.2. Comparing MAC-D machine-reading results to previous hand-coded analysis

Next, we compared the results of the LIWC coding to the previous hand-coded results for the 2436 PSF60 paragraphs. Specifically, we conducted a series of binomial logistic regressions to investigate how well the MAC-D LIWC scores predicted the previous PSF60 hand-codings.

First, demonstrating the discriminant validity of MAC-D, we found that in univariate models, each moral sub-dictionary was a significant predictor of the corresponding hand-code (for example, the Family dictionary was a significant predictor of the Family PSF60 hand-code). We also found that, in multivariate models, each moral sub-dictionary remained a significant, and was the best, predictor of the corresponding hand-code when controlling for the other six sub-dictionaries. The results of these models are summarised in Table 5, which shows the odds ratios for each predictor and each outcome variable. (For example, the Family subdictionary => Family PSF Code odds ratio of 1.82 means that for every increase of one standard deviation in Family subdictionary score, the chances of a paragraph being hand-coded Family increase 1.82 times.)

Table 5  
Odds Ratios for logistic regression of PSF Codes onto MAC-D Predictors.

		PSF Codes						
		Family	Group	Reciprocity	Heroism	Deference	Fairness	Property
MAC-D	Family	1.82	0.99	0.97	0.83	<b>1.23</b>	1.23	1.03
	Group	1.40	<b>1.98</b>	1.55	1.06	1.33	0.24	1.17
	Reciprocity	1.61	1.40	<b>3.54</b>	0.98	1.37	0.18	1.44
	Heroism	1.18	1.15	1.34	<b>1.95</b>	0.91	0.35	1.12
	Deference	1.30	0.85	0.84	1.06	<b>2.19</b>	1.27	0.87
	Fairness	1.49	1.35	0.95	0.42	0.98	<b>4.43</b>	1.20
	Property	1.13	1.21	0.99	0.25	0.67	1.85	<b>4.02</b>

Note: Bold =  $p < .001$ .



**Table 6**  
Presence of hand-coded and predicted MAC moral values across societies.

	n	Family	Group	Reciprocity	Heroism	Deference	Fairness	Property
Original PSF Hand Codes	60	78%	73%	72%	53%	58%	15%	90%
Predicted PSF Codes	60	88%	95%	88%	92%	92%	58%	85%
Predicted Non-PSF Codes	196	78%	88%	69%	78%	81%	43%	76%
All Predicted Codes	256	80%	90%	74%	81%	84%	47%	78%

Second, we dichotomised the predicted values of the univariate logistic regressions, and compared the results to the original codes to calculate inter-rater reliability. We found that the models correctly predicted the presence or absence of the moral ~80% of the time.<sup>14</sup>

However, because the proportion of true positives was low (~3%), and the proportion of true negatives was high (~97%), there was a large number of false negatives. For example, for Family, there were 138 true positives and 2298 true negatives. The model made the correct prediction 80% of the time, identifying 110/138 true positives, and 1829/2298 true negatives. But the model made the *incorrect* prediction 20% of the time, leading to 469 false positives and 28 false negatives. Hence, the true positives (110) predicted by the model were outnumbered by false positives (469); and the total positives predicted by the model (110 + 469 = 579) overestimated the actual number (138). Largely as a result of the high number of false positives, overall agreement, as measured by Cohen's kappa, was 'slight' to 'fair' (Family,  $\kappa = 0.25$ ; Group,  $\kappa = 0.10$ ; Reciprocity,  $\kappa = 0.21$ ; Heroism,  $\kappa = 0.13$ ; Deference,  $\kappa = 0.16$ ; Fairness,  $\kappa = 0.08$ ; Property,  $\kappa = 0.23$ ; all  $p$ -values < .001).

Similarly, if we aggregate these results at the society-level, we find that the models overestimate the prevalence of morals, as compared to the previously hand-coded PSF results, and hence are likely over-estimating the frequency of morals across non-PSF societies (Table 6). However, because we know that the paragraph results do not differ substantially between PSF and non-PSF societies, we can infer that the model is overestimating roughly the same across PSF and non-PSF societies, and (applying roughly the same correction) we can conclude that the prevalence of morals observed across the 60 PSF societies is approximately the same as the frequency of morals across the 196 non-PSF societies.

### 4.3. Examples

Finally, in order to get a better sense of the material, we conducted spot-checks of some of the paragraphs highlighted by the LIWC analysis.

We found many examples of the predicted moral values, as illustrated by the following tagged words in boldface (parentheticals refer to the HRAF paragraph):

- Family:** Among the Garifuna (Middle America and the Caribbean), "the highest level in the expression of **kinship** obligations" is the "**parent-child (father/mother-son/daughter)** bond". "A cardinal rule of conduct between **parents** and children is spelled out in the saying, 'Thou shall *agriaha* thine children.' The stem for *agriaha* comes from the verb meaning 'to give birth to,' 'to nourish an infant,' or 'to sustain it with life-giving support.'... As a **child** grows to adulthood, the responsibilities of *agriahouni* become reversed, with the grown son or daughter having to contribute to the support of his or her **parents**" [32].
- Group:** The Pashtun (Asia) have "a strong explicit value: that one should show neighbourliness and a positive interest in the life of all members of one's local **community**", this "unites persons most of whom are not related by kinship ties of any kind" [33].
- Reciprocity:** Among the Navajo (North America), "Another important feature of Navajo society is the view that every debt incurred must be **repaid** and that the ledger of obligations and favors received should remain in balance. ... Outside the network of kinship or even among kinsmen who are not closely related, *quid pro quo* is the rule. ... Nonrelatives usually expect to receive money, food, or other recompense for their assistance. This is considered quite normal among the Navajo, who would be loath to accept a **favor** without making some return" [34].
- Heroism:** Among the Mapuche (South America), "The chief virtue [is] physical **bravery**; then come personal **endurance**, both of pain and hardship, constancy in vengeance, and individual strength" [35]. And among Palestinians (Middle East) "**[g]enerosity** or *karam* is another old Arab virtue". "Within the village, the person who develops a reputation for being **generous** and hospitable finds that his friends increase, his word is respected, and that he is **admired** by all. On the other hand, if one gains the reputation of being stingy, one loses friends and one's prestige in the community drops" [36].
- Deference:** Among the Zulu (Africa), the "ideal ... man and woman ... has the following characteristics: (a) **Obedience** ... (b) **Respect** for his elders, or seniors ... (c) **Submissiveness** ..." [37].
- Fairness:** Among the Greeks (Europe), "The villagers take it for granted that neither **impartiality** nor altruism is possible in situations which involve one's own interests. They do not believe that anyone, even one's own brother, can act **fairly** out of principle.

<sup>14</sup> The previous paper [2] reported the correspondence between expert authors' codes and those of an independent coder. Here, the MAC-D codes are playing the role of the independent coder. The correspondence between the expert codes and the MAC-D codes demonstrates that we have succeeded in making the *implicit* knowledge of the experts *explicit* in the form of a LIWC dictionary that can operate independently.

Therefore, they feel, some mechanism is required to enforce **equity** ... the lot system effectively makes it impossible for anyone to implement favoritism” [38].

7. **Property:** Among the Orokaiva (Oceania) “[t]here is a word *igege* ... which may be translated ‘moral law’, or ... ‘moral prohibition’. The examples were such as the following: not to **steal** a man’s vegetables, canoes, spears, &c.; not to usurp his fishing rights; not to take his dog to hunt ... the mere existence of the word hints at a definite conception of right and wrong.” [39].

We also found some apparent counter-examples (though they seem to point to competition between the morals identified by MAC rather than outright devaluation of any of these morals):

**Reciprocity:** “Despite the account of Navajo reciprocity given above, another ethnographer reports, “help, whether it be in the form of labor or gifts, is thought to be ‘free’ in that it does not have to be paid for. Gifts or help are offered to create general good will, and ostensibly without the thought or expectation of **reciprocation**; and the **reciprocation**, when it does come, is not considered a return but a new act of good will. Accordingly, there is no conception of a moral obligation to **reciprocate**, as is the case in our own moral system, which often conceives of **gratitude** as a kind of indebtedness. ... Instead of saying: ‘Thanks, I hope to do the same for you some day,’ a Navaho once said to me: ‘Thanks, you can do it again some day’” [40].

**Bravery:** The (!Kung) San (Africa) “are perhaps unusual among human societies in that they attach no value to fighting. They have no ideal of **honor** or of **brave**, aggressive masculinity. There are no culturally accepted outlets for physical violence, no wrestling matches, no games of strength, and no ordeals or duels in which a man can ‘prove’ himself. The violent-tempered man and the proud man are sorry misfits among the !Kung, whose **heroes** are great curers and healers and not fighters” [41].

**Property:** Consistent with the counter-example identified by Curry et al. (2018), we found several examples of admirable theft: “to **steal** your enemy’s cattle is not **robbery** but an exploit to be proud of” (Bedouin, Middle East; [42]); “an animal **thief** not caught is the subject of admiration by all, even his victims. For it is not easy to carry off his exploit ... the pride of the unpunished **thief** and the admiration of the others for his prowess are very understandable... courage and skill, even when directed toward evil, inspire respect... [although] ... the bravery does not gainsay the illegal character of this act” (Greeks, Europe; [43]); “[c]hildren are actively encouraged to **steal**, both from other houses and from each other. Successful **theft** is considered a mark of commendable self-reliance”, and “Daring **theft** [by young men] is admired as showing courage and resourcefulness.” (Pashtun, Asia; [44]); the “**theft** of garden produce ... belonging to one village by the youths of another ... is [regarded by their relatives] as a good joke, and proof of the virility and bravery of the boys.” (Wogeo, Oceania; [45]).

We also identified examples of ‘moral molecules’ — combinations of two or more basic moral values — including filial piety (family + deference) and honour (reciprocity + heroism) [15].

“The **filial** piety demanded and expected of all Muslims ... is considered a religious duty. Ill-treatment of **parents** or **disobedience** to them constitutes a sin.” (Palestinians, Middle East; [36]).

A “revenge murder is the appropriate action in defence of **honour** after any serious insult ... any **unavenged** murder is a stain on the **honour** of the dead man’s **family**” (Pashtun, Asia; [33])

Finally, we identified situations in which the basic moral values came into conflict, creating moral dilemmas. For example, (positive) reciprocity versus heroism:

“When a Sarakatsanos receives a **favour** from another, he must show **gratitude**. ‘One **good turn deserves** another’ (ή χάρη θέλει αντίχαρη). Even if only between the acceptance of some service and its later **repayment**, **gratitude** must for a while act as a makeweight in the balance. Not to show **gratitude** for help, which has been freely given, is behaviour open to severe public reproach. But to be **grateful** is to be ‘obliged’ (παχρωμένως); and this is an admission of inequality and even of weakness. ... For neither side can admit to obligation without losing **prestige**; and both must attempt to claim the super-ordinate position in a creditor-debtor relationship.” (Greeks, Europe; [46])

And (negative) reciprocity versus group (harmony):

“the tribal moral community at large was genuinely ambivalent about a feud in its midst. On the one hand, parties who were external to the feud badly wanted the social order to be restored. On the other hand, they deeply respected a **clan** that stubbornly kept trying to even up the blood score so that it would not have to compromise its **honor**.” (Montenegrins, Europe; [47])

These examples point to the utility of MAC-D, and warrant further investigation.

## 5. Discussion

The present study used MAC-D to survey the prevalence of seven MAC morals in the ethnographic records of 256 societies – the largest-ever cross-cultural natural language analysis of moral values. Consistent with previous research, the evidence suggests that most of these morals are present in most societies, all around the world, with only minor variation across cultural region and subsistence strategy. We also found numerous examples of all seven predicted morals — as well as some counter-examples, combinations, and conflicts — in the source material.

These findings are in line with MAC’s predictions that: morality is a collection of cooperative rules; that these rules include behaviour relating to helping your family, helping your group, returning favours, acting heroically, being deferential, being fair, and respecting property; that these types of cooperation will be considered morally good; and that these moral rules will be found all

around the world.

We also found that MAC-D performed well. It correctly predicted the previous hand-codes approximately 80% of the time. We were also able to measure and to some extent diagnose the errors it made. Specifically, we were able to quantify and explain that MAC-D models overestimated the frequency of morals across paragraph and society. And crucially, the fact that MAC-D models provided similar results (and similar error) across the known PSF60 corpus and the unknown non-PSF corpus, allows us to infer that the real prevalence of morals is roughly the same across PSF and non-PSF societies. In further work, the non-PSF corpus could also be hand-coded to establish base truth and suggest revisions to MAC-D. As such, MAC-D, with its game-theoretic foundations, more comprehensive coverage of the moral domain, and validation in a global dataset rather than a parochial American/European one, has advantages over, for example, MFD and MFD 2.0. An extended MAC-D, using the methods demonstrated by Hopp et al. [30] in the development of eMFD, would likely provide an even more valuable comparison to existing approaches.

This study raises a number of questions and directions for future research. First, the present study used LIWC to estimate the cross-cultural prevalence of morals. But this method cannot be used to estimate the moral valence of cooperation. For example, it cannot distinguish between ‘helping your family is morally good’, and ‘helping your family is morally bad’, both of which would get the same score. This is probably not a major problem, given that cooperation had a positive moral valence in 99.9% cases, but it would be good to replicate this finding nonetheless. There are corpora, including well-known ones, in which family values and other elements of MAC are criticised. For instance, in the KJV translation of Matthew 10:35–6, Jesus says, “I am come to set a man at variance against his father, and the daughter against her mother, and the daughter in law against her mother in law. And a man’s foes shall be they of his own household.” Such passages are especially interesting in the context of MAC because they may indicate conflict between one moral value (in this case, Family) and another (in this case, Deference). To address potentially counter-normative uses of words associated with moral (dis)values, future research should experiment with other natural language processing tools that consider words and n-grams in context [21].

Second, the present study used MAC-D to detect seven moral values, using English words and phrases. However, the overall model fit to the previous manual codes was fairly low. And so it would be valuable to experiment with revising the words and n-grams in the dictionary, to improve the overall fit, while avoiding overfitting to this particular corpus. Future work could also use MAC-D to search more systematically for ‘moral molecules’ and ‘moral dilemmas’, by looking at those paragraphs where at least two morals are predicted to co-occur. MAC-D should also be translated into languages other than English, in order to investigate corpora in languages other than English, thereby further testing the cross-cultural robustness of MAC and the tools used to measure the elements it posits.

Third, the present study investigated the prevalence of morals in HRAF. Although the gold standard for ethnographies, this corpus has a number of limitations. HRAF was not originally collected to test MAC’s hypotheses, and what material it contains on morality was gathered somewhat serendipitously; as such, our results are likely an underestimate of the cross-cultural prevalence of these moral values. HRAF organises material by society, making it possible to compare relative frequency of morals across cultures and cultural regions. However, it is not possible using this corpus to investigate moral variation within societies — for example, between the old and the young, men and women, rich and poor, and racial and ethnic majorities and minorities [48,49]. HRAF consists of descriptions of others’ morals, as opposed to the direct expression of moral values by those who hold them. HRAF is mostly in English. For all these reasons, future research should investigate the prevalence of morals in corpora consisting of material collected more systematically, that are more representative of more local subcultural intra-community variation, that are more direct expressions of moral values by the people who hold them, in languages in addition to English. This might include: first-personal narratives, auto-ethnographies, op-eds, sermons, political speeches and party platforms, newspaper reporting, social media feeds, movie scripts, the Corpus of Contemporary American English (COCA), the Corpus of Historical American English (COHA), the Stanford Encyclopedia of Philosophy, Wikipedia, and the CHILDES database.

Fourth, the present study detected, but did not explain, minor variation between cultural regions and subsistence strategies. As such, the study did not test MAC’s prediction that moral values vary as a function of the value of different types of cooperation under different socio-economic and ecological conditions — in other words, individuals (and societies) facing different cooperative problems will prioritize different moral values [1]. Hence future research, perhaps gathering new data on the full range of moral values, using survey and questionnaire methods drawing on the Morality-as-Cooperation Questionnaire (MAC-Q) developed by Curry et al. [16], from representative cross-cultural samples, will be needed to investigate the relationship between particular moral values and relevant individual- and societal-level indicators of cooperation — such as family size and dispersal, group size, mobility, subsistence strategy, reliance on trade, frequency of warfare, degree of inequality, political structure, age structure, resource base, and territory size.

Finally, a limitation of the paper is that, as we started this investigation in 2020, we aimed to use standard methods available at the time, favoring a transparent, expert-guided (yet more traditional) approach, as opposed to more ‘black-boxed’ techniques such as deep neural networks. More advanced methods that could complement our work were not available. Obviously, things have moved on since then, with ChatGPT and other language models derived from Bidirectional Encoder Representations from Transformers (BERT); and future work should employ these methods.

This study introduces to the field MAC-D, a new dictionary for surveying moral corpora. Future work developing the theory and revising the tool will further our understanding of the nature and content of human morality.

## Significance statement

The empirical study of morality has hitherto been conducted primarily in WEIRD contexts and with living participants. This paper addresses both of these shortcomings by examining the global anthropological record. In addition, we develop a novel methodological tool, the morality-as-cooperation dictionary, which makes it possible to use natural language processing to extract a moral signal from

text. We find compelling evidence that the seven moral elements posited by the morality-as-cooperation hypothesis are documented in the anthropological record in all regions of the world and among all subsistence strategies. Furthermore, differences in moral emphasis between different types of cultures tend to be non-significant and small when significant. This is evidence for moral universalism.

## Funding

The research leading to this paper was supported by The Toyota Foundation (TYTID: D17-R-0506). Oliver Scott Curry's time working on this project was supported by [Kindness.org](https://www.kindness.org).

## Conflict of interest disclosure

The authors declare no conflicts of interest.

## Ethics declarations

Review and/or approval by an ethics committee was not needed for this study because it does not involve human participants and is based on existing published material.

## CRedit authorship contribution statement

**Mark Alfano:** Writing – review & editing, Writing – original draft, Visualization, Software, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Marc Cheong:** Writing – review & editing, Writing – original draft, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Oliver Scott Curry:** Writing – review & editing, Writing – original draft, Visualization, Validation, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Oliver Scott Curry reports financial support was provided by Toyota Foundation. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgements

We are grateful to Carol Ember and Michael Fischer for providing access to the HRAF material, to James Wilkinson, Michal Misiak, Darragh Hare, Christine Pelican, and Simon Coghlan for helping us crowd-source the dictionaries developed in this paper, and to Christine Pelican for assistance with data analysis.

## Appendix A Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.heliyon.2024.e25940>.

## References

- [1] O.S. Curry, Morality as cooperation: a problem-centred approach, in: T.K. Shackelford, R.D. Hansen (Eds.), *Evol. Moral.*, Springer International Publishing, 2016, pp. 27–51.
- [2] O. Curry, D. Mullins, H. Whitehouse, Is it good to cooperate? Testing the theory of morality-as-cooperation in 60 societies, *Curr. Anthropol.* 60 (2019) 47–69.
- [3] D. Lieberman, J. Tooby, L. Cosmides, Does morality have a biological basis? An empirical test of the factors governing moral sentiments relating to incest, *Proc. R. Soc. Lond. Ser. B Biol. Sci.* 270 (2003) 819–826, <https://doi.org/10.1098/Rspb.2002.2290>.
- [4] W.D. Hamilton, The genetical evolution of social behaviour, *J. Theor. Biol.* 7 (1–16) (1964), [https://doi.org/10.1016/0022-5193\(64\)90039-6](https://doi.org/10.1016/0022-5193(64)90039-6), 17–52.
- [5] D.K. Lewis, *Convention: a Philosophical Study*, Harvard University Press, Cambridge, Massachusetts, 1969.
- [6] J. Royce, *The Philosophy of Loyalty*, Macmillan, New York, 1908.
- [7] L. Cosmides, J. Tooby, Social exchange: the evolutionary design of a neurocognitive system, in: M.S. Gazzaniga (Ed.), *New Cogn. Neurosci.*, MIT Press, Cambridge, MA, 2005, pp. 1295–1308.
- [8] R.L. Trivers, The evolution of reciprocal altruism, *Q. Rev. Biol.* 46 (1971) 35–57, <https://doi.org/10.1086/406755>.
- [9] O.S. Curry, The conflict-resolution theory of virtue, in: W.P. Sinnott-Armstrong (Ed.), *Moral Psychol.*, MIT Press, Cambridge, Massachusetts, 2007, pp. 251–261.
- [10] D. Sznycer, L. Al-Shawaf, Y. Bereby-Meyer, O.S. Curry, D. Smet, E. Ermer, J. Tooby, Cross-cultural regularities in the cognitive architecture of pride, *Proc. Natl. Acad. Sci. USA* 114 (2017) 1874, <https://doi.org/10.1073/pnas.1614389114>.
- [11] D.M. Messick, Equality as a decision rule, in: B. Mellers, J. Baron (Eds.), *Psychol. Perspect. Justice*, Cambridge University Press, Cambridge, 1993, pp. 11–31.
- [12] B. Skyrms, *Evolution of the Social Contract*, Cambridge University Press, Cambridge, 1996.
- [13] H. Gintis, The evolution of private property, *J. Econ. Behav. Organ.* 64 (2007) 1–16, <https://doi.org/10.1016/j.jebo.2006.02.002>.
- [14] J. Graham, J. Haidt, B.A. Nosek, Liberals and conservatives rely on different sets of moral foundations, *J. Pers. Soc. Psychol.* 96 (2009) 1029.

- [15] O.S. Curry, M. Alfano, M.J. Brandt, C. Pelican, Moral molecules: morality as a combinatorial system, *Rev. Philos. Psychol.* 13 (2022) 1039–1058, <https://doi.org/10.1007/s13164-021-00540-x>.
- [16] O. Curry, M. Chesters, C. Lissa, Mapping morality with a compass: testing the theory of ‘morality-as-cooperation’ with a new questionnaire, *J. Res. Pers.* 78 (2019) 106–124.
- [17] M. Zakharin, O.S. Curry, Others in QIMR team, N.G. Martin, G.J. Lewis, T.C. Bates, Modular Morals: The Genetic Architecture of Morality as Cooperation, (n.d.), <https://osf.io/kms9c/>.
- [18] M. Buyukozer Dawkins, F. Ting, M. Stavans, R. Baillargeon, Early moral cognition: a principle-based approach, *Cogn. Neurosci.* VI, (2019) 7–16.
- [19] C.S. Otterbein, Methods in cross-cultural psychology, *Am. Psychol.* 33 (1978) 857.
- [20] C.R. Ember, M. Ember, *Cross-Cultural Research Methods*, second ed., Altamira Press, 2009.
- [21] J. Pennebaker, R. Boyd, K. Jordan, K. Blackburn, *The Development and Psychometric Properties of LIWC2015*, University of Texas at Austin, 2015, <https://doi.org/10.15781/T29G6Z>.
- [22] J. Pennebaker, *The Secret Life of Pronouns: what Our Words Say about Us*, Bloomsbury Press, 2011.
- [23] R.L. Boyd, H.A. Schwartz, Natural language analysis and the psychology of verbal behavior: the past, present, and future states of the field, *J. Lang. Soc. Psychol.* 40 (2021) 21–41.
- [24] E. Abedin, M. Ferreira, R. Reimann, M. Cheong, I. Grossmann, M. Alfano, Exploring intellectual humility through the lens of artificial intelligence: top terms, features and a predictive model, *Acta Psychol.* 238 (2023) 103979.
- [25] M. Alfano, A. Higgins, J. Levernier, Identifying virtues and values through obituary data-mining, *J. Value Inq.* 52 (2018) 59–79.
- [26] J. Hoover, G. Portillo-Wightman, L. Yeh, S. Havaladar, A.M. Davani, Y. Lin, M. Dehghani, Moral Foundations Twitter Corpus: a collection of 35k tweets annotated for moral sentiment, *Soc. Psychol. Personal. Sci.* 11 (2020) 1057–1071.
- [27] J. Trager, A.S. Ziabari, A.M. Davani, P. Golazazian, F. Karimi-Malekabadi, A. Omrani, M. Dehghani, *The Moral Foundations Reddit Corpus*, 2022.
- [28] W.E. Padfield, E.M. Buchanan, *Moral Foundations of U.S. Political News Organizations*, 2020. <https://osf.io/k4cm8/>.
- [29] N. Buttrick, R. Moulder, S. Oishi, Historical change in the moral foundations of political persuasion, *Pers. Soc. Psychol. Bull.* 46 (2020) 1523–1537.
- [30] F.R. Hopp, J.T. Fisher, D. Cornell, R. Huskey, R. Weber, The extended Moral Foundations Dictionary (eMFD): development and applications of a crowd-sourced approach to extracting moral intuitions from text, *Behav. Res. Methods* 53 (2021) 232–246.
- [31] M. Gamer, J. Lemon, I. Fellows, P. Singh, Package ‘irr’, 2019. <https://cran.r-project.org/web/packages/irr/irr.pdf>. (Accessed 29 July 2020).
- [32] J.O. Palacio, Kin ties, food and remittances in a Garifuna village in southern Belize, *Diet Domest. Life Soc.* 119–146 (1991) 263–272. <https://ehrafworldcultures.yale.edu/document?id=sa12-022>.
- [33] F. Barth, Political leadership among swat pathans, *Monogr. Soc. Anthropol.* 8 (1965) 143. <https://ehrafworldcultures.yale.edu/document?id=au04-004>.
- [34] J.F. Downs, Navajo, *Case Stud. Cult. Anthropol.* 1972, p. 136, vol. i, <https://ehrafworldcultures.yale.edu/document?id=nt13-185>.
- [35] R.E. Latham, *Ethnology of the araucanos*, *J. R. Inst. G. B. Irel.* 39 (1909) 334–370, 2 plates.
- [36] A.M. Lutfiyya, Baytin a Jordanian village: a study of social institutions and social change in a folk community, *Stud. Soc. Anthropol.* v. 1 (1966) 202. <https://ehrafworldcultures.yale.edu/document?id=m013-038>.
- [37] O.F. Raum, *The Social Functions of Avoidances and Taboos Among the Zulu*, *Monogr. Zur Völkerkd.* 1973, p. 576, xiv, <https://ehrafworldcultures.yale.edu/document?id=fx20-015>.
- [38] E. Friedl, Vasilika: a village in modern Greece, *Case Stud. Cult. Anthropol.* 18 (1963) 110. <https://ehrafworldcultures.yale.edu/document?id=eh01-002>.
- [39] E.E. Williams, H. Murray, Orokaiva Society, 1930. <https://ehrafworldcultures.yale.edu/document?id=oj23-001>.
- [40] J. Ladd, A systematic reconstruction of Navaho ethics, *Struct. Moral Code Philos. Anal. Ethical Discourse Appl. Ethics Navaho Indians.* xv (1957) 474. <https://ehrafworldcultures.yale.edu/document?id=nt13-196>.
- [41] R.B. Lee, *!Kung Bushmen of Botswana*, *Hunt. Gatherers Today*, 1972, pp. 327–368. <https://ehrafworldcultures.yale.edu/document?id=fx10-040>.
- [42] H.R.P. Dickson, Arab of the desert: a glimpse into Badawin life in Kuwait and Sau’i Arabia, 1951. <https://ehrafworldcultures.yale.edu/document?id=mj04-001>.
- [43] G.V. Kavvadias, F. Schütze, Mediterranean pastoral nomads: the Sarakatsani of Greece, *Collect. Sci. Hum. Appliquées.* (1965). HRAF MS: ii, 498 [original: x, 444 ], <https://ehrafworldcultures.yale.edu/document?id=eh01-013>.
- [44] C. Lindholm, Generosity and Jealousy: the Swat Pukhtun of Northern Pakistan, 1982. <https://ehrafworldcultures.yale.edu/document?id=au04-003>.
- [45] H.I. Hogbin, *Social reaction to crime: law and morals in the Schouten Islands, New Guinea*, *R. Anthropol. Inst. G. B. Irel.* 68 (1938) 223–262, plates X–XII.
- [46] J.K. Campbell, *Honour, Family and Patronage: a Study of Institutions and Moral Values in a Greek Mountain Community*, 1964. <https://ehrafworldcultures.yale.edu/document?id=eh01-011>.
- [47] C. Boehm, *Blood Revenge: the Anthropology of Feuding in Montenegro and Other Tribal Societies*, 1984. <https://ehrafworldcultures.yale.edu/document?id=ef05-002>.
- [48] M. Moody-Adams, *Fieldwork in Familiar Places: Morality, Culture, and Philosophy*, Harvard University Press, 1997.
- [49] M. Alfano, *Moral Psychology, An Introduction*, Polity, 2016.
- [50] M. Dehghani, R. Boyd, *Handbook of Language Analysis in Psychology*, Guilford Press, 2022.
- [51] K. Benoit, K. Watanabe, H. Wang, P. Nulty, A. Obeng, S. Mueller, A. Matsuo, J. Kuha, W. Lowe, C. Mueller, L. Young, S. Soroka, I. Fellows, *Quanteda: Quantitative Analysis of Textual Data*, 2020. <https://cran.r-project.org/web/packages/quanteda/index.html>. (Accessed 29 July 2020).
- [52] G. Nicolas, X. Bai, S.T. Fiske, Comprehensive stereotype content dictionaries using a semi-automated method, *Eur. J. Soc. Psychol.* 51 (2021) 178–196.
- [53] P. Gloor, A.F. Colladon, F. Grippa, Measuring ethical behavior with AI and natural language processing to assess business success, *Sci. Rep.* 12 (2022) 10228.
- [54] L. Bulla, S.D. Giorgis, A. Gangemi, L. Marinucci, M. Mongiò, Detection of morality in tweets based on the moral foundation theory, in: *Int. Conf. Mach. Learn. Optim. Data Sci.*, Springer Nature Switzerland, Cham, 2022, pp. 1–13.
- [55] G.A. Miller, *WordNet: a lexical database for English*, *Commun. ACM* 38 (1995) 39–41.
- [56] S. Bird, E. Loper, E. Klein, *Natural Language Processing with Python*, O’Reilly Media Inc, 2009.
- [57] K. Atkinson, *SCOWL (And Friends)*, 2009. <http://wordlist.aspell.net/>. (Accessed 31 July 2020).