Modeling the Sacred: Considerations when Using Religious Data in Natural Language Processing

Anonymous ACL submission

Abstract

This position paper concerns the use of religious data in Natural Language Processing (NLP), which is of special interest to the Ethics of NLP. Religious texts are expressions of values and cultural practices that relate to deeply held convictions, and machine learned models have a propensity to reproduce cultural values and biases encoded in their training data. Furthermore, translations of religious texts are increasingly being used by researchers and developers, especially when language data is scarce. This repurposes the translations from their original uses and motivations, which often involves attracting new followers. With these in mind, this paper discusses the considerations of using religious texts for the development of language technologies, including concerns around cultural rights.

1 Introduction

The Association for Computational Linguistics (ACL) is a secular institution. Its constitution, resolutions and policies make no mention of religion other than forbidding harassment on the basis of religion.[1] Nevertheless the Christian Bible and the Islamic Quran[2] are increasingly being used in the scientific and professional activities of ACL, as measured by papers published in the ACL Anthology (Figure 1). Some of the reasons that NLP researchers use the Bible are aptly expressed by Resnik et al. (1999). The Bible is the world's most translated book, with translations in over 2,000 languages, and often multiple translations per language. Furthermore, great care is taken with the translations, so from an NLP perspective data quality is high. It is often easily available in electronic form, and is in the public domain, hence free to use. It has a standard structure which allows parallel alignment verse-by-verse. (And some of these reasons, or similar ones, hold of the Quran too.) For these reasons, as recently as 2006 it was said to be "perhaps surprising that the Bible has not been more widely used as a multilingual corpus by the computational linguistics and information retrieval community" (Chew et al., 2006).

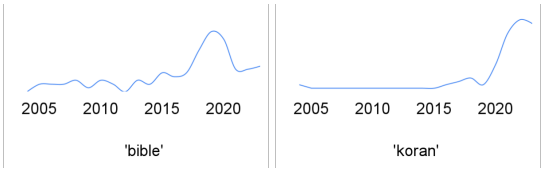

Figure 1: Trend sparklines for counts of papers in the ACL Anthology mentioning 'bible' or 'koran' (see §3).

Despite the increasing use of sacred texts in NLP, the ethical considerations around such use does seem not to have received prior attention (other than a brief mention by Mager et al. (2023)). This position paper contends that responsible secularism demands engaging with the ethical considerations of the use of sacred texts. If one should never speak of religion in polite company, then perhaps ACL forums should be less polite. In Section 2, we provide relevant background, and summarize a debate within the field of academic linguistics concerning its disciplinary relationship with missionary linguistics. In Section 3, we present a study of papers in the ACL Anthology mentioning sacred texts, as well as four recent case studies that illustrate ways in which NLP research uses sacred texts. Section 4 discusses a range of ethical considerations when using sacred texts in NLP. In doing so, we consider a range of approaches to the topic, including ethical theories, Indigenous perspectives,[3] human rights, and the AI principles commonly espoused by institutions. Our goal in doing so is not to eval-

[1] www.aclweb.org, accessed September 2023

[2] This paper follows several style guides in using "Quran", although mentions of the alternate Latinization "Koran" are also considered in the corpus studies we report on.

[3] Following several style guides, we capitalize the first letter of "Indigenous". See, e.g., https://www.sapiens.org/language/capitalize-indigenous/.

Religion	Sacred texts	Est. 2020 population	Proselytizing
Christianity	Bible, New Testament, Old Testament	2,382,750,000	Yes
Islam	Quran (alt. spellings include Koran), Old Testament	1,907,110,000	Yes
Hinduism	Vedas, Upanishads, Puranas	1,161,440,000	No
Buddhism	Tripitaka, Mahayana Sutras, Tibetan Book of the Dead	506,990,000	Yes
Traditional Chinese Religion	Zhuangzi, Tao-te Ching, Daozang	310,000,000	No
Judaism	Talmod, Torah, Tanakh, Old Testament	14,660,000	No

Table 1: Some major world religions and their texts. Population estimates are from the US-based Pew Research Center (www.pewresearch.org), which conducts demographic and other research.

uate past NLP projects that use religious texts, but rather to encourage more reflecting in and on future work. Based on these considerations, we then make some recommendations for the NLP community in Section 5, concerning cultural standpoints, cultural knowledge gaps, and power dynamics between global and marginalized cultures.

2 Background

2.1 Religion

Precisely defining what constitutes a religion might be notoriously difficult (see e.g., Spiro, 2013; Neville, 2018), and lies beyond the scope of this paper. Common properties of religions center around giving meaning to existence, and include: a) moral values concerning which actions are right or wrong, b) spiritual beliefs, including what happens after death, c) theistic beliefs about gods, or spirit beings, d) rituals around birth, initiation, obtaining adulthood, marriage, and death, e) stories and mythologies concerning topics such as the origin of the world, the origin of humans, etc, f) kinship systems and marriage practices, g) artistic practices, including songs, dances, and visual arts, h) significant locations, including buildings, sites, and homelands; and in some cases i) a language which plays a special role. These are all closely related to questions concerning values: moral values, spiritual values, cultural values, aesthetic values, historic values, and even linguistic values.

Like languages, religions may exhibit regional variation and incorporate local practices, so thinking of them as discrete entities may be somewhat misleading. Some widely cited estimates put the number of worldwide religions at several thousand, although these claims are disputed. What seems more certain is that the imminent extinction of many Indigenous languages will accompany an "impending loss of so many religions and worldviews" (Harrison, 2007, p. 153). Acknowledging these challenges, we nevertheless provide a summary of some of the world's most populous religions in Table 1. Surveying or defining each of these religions is beyond the scope of this paper. Also not included here are the various sects and branches within each religion (e.g., Catholicism), nor texts which might be important only to specific branches (e.g., *The Book of Mormon*).

One important distinction is that between proselytizing and non-proselytizing religions. The former attempt to convert new populations, whereas the latter do not. The former are more intricately related to historical practices of colonialism—especially in Africa, the Americas, Asia, and the Pacific—and hence also to neocolonial legacies. Some religions hold that a certain language is privileged for communicating sacred texts to the faithful, while on the other hand Protestant Christianity exemplifies a commitment to communicating in vernacular languages. (Article XXIV of the Articles of Religion of the Anglican Church calls for "such a Tongue as the people understandeth".)

2.2 The Academy and Bible Translation

The September 2009 issue of the journal *Language* has a special feature of 5 articles by anthropologists and linguists concerning the relationships between the US-based Bible translation organization SIL International (SIL) and academic linguistics. In this issue, Dobrin and Good (2009) explore how academic linguists have at times become reliant on, and benefitted from, the technological infrastructures of SIL, in part because creating and maintaining these infrastructures has not been valued by the academy. This "partnership of convenience" causes tensions between differing objectives, and raises questions about what kind of relationships secular research institutions should have with organizations with very different agendas. These practices presage similar ways in which some areas of NLP research have become reliant on Bible translations. Many linguists and NLP practition-

ers working on Indigenous languages see their research as addressing issues of human rights and cultural extinction. However, as Dobrin and Good point out, languages which are most endangered are least likely to receive SIL's attention. Handman (2009) draws attention to how SIL ideology separates linguistic identity from religious identity, differing from UNESCO's position that sustaining endangered languages entails sustaining cultural worldviews, knowledge systems, and identity practices. Epps and Ladley (2009) argue that evangelical success entails the displacement or transformation of traditional beliefs, often leading to social upheaval, and argues that the academy has a moral interest in supporting local self-determination which is at odds with evangelical agendas.

3 The Use of Sacred Texts in NLP

In this section we demonstrate that the use of religious data is common in the field of NLP when using machine learning to train models. We focus primarily on the publications of the NLP research community, as represented by the searchable Anthology of the Association of Computational Linguistics (a.k.a., the ACL Anthology) (Bird et al., 2008).[4] We acknowledge that this corpus might not be representative of the entire field of NLP, e.g., missing relevant work such as (Chandra and Ranjan, 2022; Bashir et al., 2023), and NLP projects in industry might not be well represented in the ACL Anthology. Another limitation is that we exclude publications in languages other than English.

3.1 Sacred Texts in the ACL Anthology

The number of ACL Anthology entries for each of the texts listed in Table 1 are shown in Table 2. Thousands of papers in the ACL Anthology seem to use religious texts. There is a strong bias towards the texts of the monotheistic Abrahamic religions of Judaism, Christianity, and Islam. All three originate in the Middle East, with original versions of their texts in Hebrew, Ancient Greek, and Classical Arabic. However Christian texts have been widely distributed in European languages following the Protestant Reformation, and the global spread of Christian influence is associated with European history and the colonial and missionary practices of European cultures.

To understand *when* sacred texts have been used by the NLP community, we examined the year of publication of the first 100 results (sorting by relevance, not by date) on the ACL Anthology for each of the terms 'bible' and 'koran'. (The ACL Anthology search interface only made these available.) We omitted 12 'bible' search results from our analysis due to either: not being in English (2); not being research papers (6, e.g., book reviews, invited talks, or proceedings); or being duplicates (4). We omitted 36 'koran' search results from our analysis due to either: not being research papers (16); being false positive search results (13, e.g., 10 had typos for 'Korean'); or only using the word 'koran' in references (3), footnotes (2) or content generated by a model (2). We manually coded each paper for its year of publication, NLP application domain and which languages it covered. 150 of the resulting 152 papers spanned from 2004 to July 2023, and of these we see increasing use of the Bible and the Koran in NLP research over time (see Figure 1 in Section 1[5]), including over 60% of papers being published between 2019 and 2023 alone.

To understand *how* sacred texts are used in NLP, we analyzed the 88 papers mentioning the Bible described above. A variety of application domains were represented, including sentiment analysis, named entity tasks, CLIR, patronizing language detection, and various morpho-syntactic analysis tasks. However the most common applica-

Search term/phrase	Results [min, max]
bible	[1920, 3890]]
quran	[291, 547]
new testament	[294, 294]
koran	[131, 248]
old testament	[73, 206]
torah	[25, 153]
talmud	[21, 22]
vedas	[22, 51]
tripitaka	[7, 7]
upanishads	[6, 6]
mahayana sutras	[4, 4]
tanakh	[3, 4]
zhuangzi	[3, 3]
puranas	[3, 3]
tibetan book of the dead	[0, 0]
tao-te ching	[0, 0]
daozang	[0, 0]

Table 2: Number of search results for religious texts in the ACL Anthology on August 10, 2023. Since ACL Anthology search result counts are non-deterministic, we report the min and max of 10 searches for each term.

[4] www.aclanthology.org, accessed August 2023.

[5] Since our 2023 sample was limited to Jan-July, paper counts for 2023 are multiplied by 12/7 when creating sparklines, to make them comparable with previous years.

tion domain was machine translation (22%), while many (18%) introduced a new corpus or lexical resource. Three papers were concerned specifically with Bible translation, four with literary analysis of the Bible, and five with language modeling or pretraining. There seems to be a recent trend towards papers handling very large numbers of languages, with ten papers since 2014 in our sample handling over 500 languages. 48% of the 88 papers concerned one or more Indigenous languages.[6] The 64 ACL papers mentioning 'koran' were less varied and overwhelmingly used verses included in machine translation evaluation datasets.

3.2 Four Contemporary Case Studies

To complement the broad analysis of the previous section, we also report here in more detail on some recent noteworthy NLP papers. These illustrate in more detail ways in which the NLP community is encountering and using religious texts.

Our first example is a paper which was uploaded to arXiv, a popular online archive for computer science papers, in May 2023 (Pratap et al., 2023). This work aims to improve Speech Recognition and Text To Speech Synthesis for over a thousand languages. They train their model using translations of the New Testament, as well as audio of readings of those translations, obtained from Faith Comes by Hearing (faithcomesbyhearing.com), goto.bible and bible.com. They also use spoken recordings in many languages, without paired texts, of Bible stories, evangelistic messages, scripture readings, and songs, obtained from Global Recordings Network (globalrecordings.net), whose mission is to communicate "the Good News of Jesus Christ" via a strategy of recording, distribution, and promotion.

Our second example was awarded the ACL Area Chair Award for best Multilingualism and Cross-Lingual NLP paper in July 2023 (ImaniGooghari et al., 2023). This work aims to scale language models to 500 languages. They "crawl or download" data from 150 sources, including religious texts and observe a "higher proportion of religious data" compared to previous comparable work. Parallel verses from Bible translations are used for model training and testing, and performance is reported for Sentence Retrieval from the Bible.

Our third example concerns JW300 (Agić and Vulić, 2019), a dataset of around 100k sentences in each of 300 languages crawled from jw.org, a website run by the US-based Jehovah's Witnesses, a Christian denomination. A majority of the texts come from the Jehovah's Witnesses' magazines *Awake!* and *Watchtower*. Released as a corpus in 2019, JW300 has been cited over 150 times as of August 2023. The African grassroots open-source NLP project MASAKHANE (masakhane.io) had been using JW300 to train Machine Translation models, until receiving legal advice in 2023 that this was breaching copyright. A subsequent request by MASAKHANE to the Jehovah's Witnesses for permission to use the data was declined.[7]

Our final recent example concerns the release of MADLAD-400, a new text dataset containing 3T tokens in 419 languages (Kudugunta et al., 2023). It uses 2022 snapshots of the CommonCrawl web crawl (commoncrawl.org) and the paper was uploaded to arXiv in September 2023. Auditing of a preliminary version of the dataset, spanning 498 languages, revealed that for 141 languages there were "significant amounts" of Bible data. Significant amounts of Jehovah's Witnesses data was also found for 37 languages, and of Church of Jesus Christ of the Latter Day Saints (LDS) data for 2 languages. (No Quran data was reported to be found in significant amounts.)

4 Considerations

Having demonstrated above that religious texts have been used in thousands of NLP papers, we now discuss some of the ethical considerations. Our goal is not to critique the work discussed in the previous section, but rather to provide a toolbox for assisting critical thinking in the future. Following a call for NLP researchers to focus their ethical considerations on power relations between technologists and communities (Blodgett et al., 2020)—and cogniscant that we are in the International Decade of Indigenous Languages (2022–2032)—we give special attention to relationships between local Indigenous communities and global projects, for which large power disparities exist and for which "the asymmetry of power is the cause of domination" (Mager et al., 2023).

[6] These categorizations were done by the author, taking into account historical and social context, however an ideal approach might engage with language communities to understand whether they consider themselves Indigenous.

[7] See, e.g., https://walledculture.org/a-blatant-no-from-a-copyright-holder-stops-vital-ling or https://www.youtube.com/watch?v=mbkuRZkg1RY.

4.1 On Cultural Relativism

One possible objection to the claim that such considerations are needed is based on cultural relativism. Such an objection would argue that by using sacred texts for an extended period, the ACL community has demonstrated its acceptance of such practices. That is, such practices should be judged as acceptable by the norms of the ACL community.

We counter that such an objection would be stronger if the ACL community both had a stronger history of reflexive practices, and was more culturally diverse. Compared to many other disciplines, we find that ACL's interest in ethics to be relatively recent. The first Workshop on Ethics in NLP was held in 2017 (Hovy et al., 2017). "Ethics and NLP" was not a possible submission topic for the ACL conference until 2020. The ACL adopts the ACM Code of Ethics,[8] a general code for computing professionals which makes no mention of working with cultural data such as language. In 2022, a Responsible Research Checklist was introduced. In 2023, seemingly for the first time, the ACL encourages (but does not require) an ethics statement. Unlike some other disciplines, positionality statements are rare in ACL papers, with "researcher positionality" and "author positionality" each having only a single (and recent) result in the ACL Anthology as of September 2023. We also find that the ACL community does not represent the diversity of the world's local languages and religions; in fact disparities in sources of ACL publications might be increasing (Rungta et al., 2022).

4.2 The Veil of Religious Ignorance

A useful starting premise may be that ethical consideration of the papers in the journals and proceedings of the ACL community should not be biased against nor towards any religions. Rawls' Veil of Ignorance suggests a guide (Rawls, 1971): that one's consideration of the use of sacred texts within NLP should proceed as if one were in ignorance which religion (if any) we each belong to. For example, how would I feel about NLP's reliance on texts from major religions if my own culture and religion might be marginalized and endangered?

4.3 Etic NLP and Emic NLP

Although the Veil of Ignorance might seem useful in theory, our ability in practice to truly avoid being informed by our own cultural backgrounds and affiliations is highly questionable. We must instead accept that we have cultural standpoints. The emic/etic distinction originated in linguistics in the 1950s for describing different standpoints for language research (Mostowlansky and Rota, 2020). *Emic* is commonly used to describe research on a culture from the perspective of people of that culture. This contrasts with *etic* research, which takes an outsiders perspective. We propose it is useful to talk of *etic NLP* and *emic NLP*, according to whether the language technology is for our own linguistic cultures or those of others. Similarly, when NLP handles religious texts, we can distinguish research problems and applications which are within the researcher's own religious context, from those applications which impact those having other religious beliefs (e.g., translation for the purpose of proselytizing).

The more different that a culture is from our own, the less we should expect to understand the ethical considerations that those of that culture may have around language technologies. Of special consideration here are Indigenous cultures, which perhaps diverge the greatest from the majority global cultures. These are grappling with historic and current marginalization which often places their worldviews, languages and religions at risk of extinction, and their voices are among the least likely to be represented in ACL's prestigious conference presentations, plenaries, panels, and journals.

4.4 A Range of Lenses

In the previous subsections, we have already begun to lay the groundwork for our position that the ACL community should adopt an acknowledge our own individual cultural standpoints, ignorances of other cultures, and relationships to global and marginalized cultures. With this in mind, we now consider the use of sacred texts in NLP from a range of moral and sociological lenses.

The term **consequentialism** refers to a family of normative theories which emphasize the importance of considering the consequences of actions. An example is **utilitarianism**, which holds that the best action is one that maximizes wellbeing and minimizes suffering. It has been argued that AI research often implicitly adopts a utilitarian lens (e.g., Hutchinson et al., 2022). One challenge with applying these theories in an NLP research context is that in practice research activities are often far

[8] https://www.aclweb.org/portal/content/acl-code-ethics

removed from applications, and any eventual path between the two can be unknowable or uncertain. Even if the impacts on users of NLP applications can be known, challenges arise in calculating aggregate benefits across disparate stakeholders with different objectives. Perhaps the easiest outcomes to reason about are the impacts of NLP research and applications on the lives of the researchers and developers themselves, since research papers and software bring rewards from within the NLP community itself, in the forms of kudos, citations, career progression, etc. How can these be weighed against the possible risks of dignitive harms, e.g., if believers feel that NLP systems operating with a non-trivial error rate are offensively trivialising sacred texts? If copyrighted parallel texts are purchased from a proselytizing organization, what are the downstream impacts of the organization re-investing those monies into efforts to convert?

The term **deontology** refers to a normative theory which posits that there are rules or principles which determine the rightness of wrongness of actions, rather than the consequences. Within an NLP context, this might lead to a focus of on the upstream actions such as sourcing of data, rather than downstream actions such as usage of NLP applications. Within the context of sacred texts and their translations, this might lead to questions such as: Were the right people involved in the creation of a dataset of sacred texts? Did they have the right roles and relationships from within the perspective of followers of the religion? Do translators of sacred texts have the right specialist translation skills and cultural knowledge? Were the translators paid fairly? Generally, was the dataset collected in a manner aligned with best research practices, e.g., as operationalized by research ethics boards? Prabhumoye et al. (2020) discuss the importance of informed consent for deontological approaches to NLP ethics, and community-level consent might be an appropriate lens for thinking about communities of religious practice.

We use the term **AI Principles** here to refer broadly to the hundreds of sets of principles for responsible and ethical AI that have been released by companies and governments in recent years. Although all unique in their own way, the sets of principles also have many facets in common (Floridi and Cowls, 2022), some of which we discuss here. From a *safety* perspective, could the use of sacred texts be used in misinformation or disinformation? For example, could an NLP system trained on sacred texts be used in such as way that, deliberately or accidentally, leads to false beliefs about a religion? From a *privacy* perspective, when dealing with speech recordings of sacred texts, does the inherently identifiability of human voices introduce concerns which can't be mitigated? From a *bias and fairness* perspective, what kinds of cultural biases and encodings of values are present in the sacred texts and likely to be reproduced by systems trained on the those texts? Given that sacred texts are poor representations of other linguistic domains (see discussion in e.g., Mayhew et al., 2017; Adelani et al., 2021), what kinds of system biases are likely to result when systems trained on sacred texts are used in other domains? From an *accountability* perspective, who is accountable for inappropriate behaviours of systems trained on religious texts? What are their accountabilities are there to followers of religions if systems produce offensive religious language or misinformation? From a *transparency* perspective, have datasets incorporating sacred texts followed best practices around dataset documentation (e.g. Bender and Friedman, 2018; Gebru et al., 2021; Pushkarna et al., 2022)?

One topic in the **sociology of technology** is concerned with how power is distributed. As Winner (2017) has argued, artefacts are political in the sense that they make it easier for some people to do some things, and perhaps more difficult for other people to do other things. Languages, like language technologies, have politics. Following calls by Blodgett et al. (2020), we might focus our considerations on which NLP tools trained on sacred texts serve to re-arrange power, by considering which actions are encouraged or discouraged and by who? And we can consider the etic/emic distinction (§4.3): how are those within or outside the religious community empowered or disempowered, compared to others? For example, accurate machine translation into a language might inhibit the need for paying human translators from a marginalized community.

As discussed in Section 3, many NLP papers using the Bible use translations of its texts into Indigenous languages. There are many **Indigenous lenses** on responsible research, technology and data practices, as well as guidelines for non-Indigenous researchers working with Indigenous data (e.g., Carroll et al., 2020; National Health and Medical Research Council (Australia), 2018;

Taiuru, 2021). However, despite the plurality, some common themes emerge regarding concerns when working with Indigenous communities and their cultural data, including linguistic data. One theme is around *relationships*: Does the researcher have a relationship with the Indigenous community which involves generosity, reciprocity, humility, responsibility, obligations, and care? Another is around *benefits*: Does the Indigenous community benefit, collectively and individually, from the NLP project? Are benefits shared equitably? Is the community empowered by the project, and are their capabilities and capacities improved? Are possible harms appropriately mitigated? A third theme is around *culture*: Does the project respect the spirit and integrity for all facets of the Indigenous culture? Are cultural concerns around secrecy and privacy respected? Does the NLP project help to maintain Indigenous culture and connections to culture, assisting with continuity and mitigating threats to extinction? A fourth theme is around *control*: With concerns around Indigenous data sovereignty in mind, are Indigenous people in control of their cultural data? Can projects and their outcomes be contested by the Indigenous community?

Further considerations around **international laws and human rights** concern NLP's use of sacred texts translated into Indigenous languages. As tools of colonizing projects, such translations have been described as a "well documented example of the non-ethical misuse of translation" (Mager et al., 2023). Kenyan human rights scholar Makua Mutua describes at length what he calls the "basic contradictions" between proselytizing religions and Indigenous cultures (Mutua, 2004). Observing that religion is woven into every aspect of social and cultural life in Indigenous cultures, including dances, ceremonies, rites, and marriage practices, Professor Mutua argues that the meeting of such cultures with proselytizing Christian and Islam faiths amounts to "cultural genocide". In some cases, this characterization seems valid, and also true of Indigenous encounters with other dominating ideologies, for example the brutal repression of Indigenous religions by Soviet Russia led to extermination of religions without a trace (Harrison, 2007, p. 152). In other cases, "cultural genocide" may be too strong a phrase for the complex realities of how local cultures respond syncretically to proselytizing cultures. For example, Australian historian Laura Rademaker describes how on Groote Eylandt, an island in remote Northern Australia, the Indigenous Anindilyakwa people reinterpreted Christianity of the missions in their own ways leading to a "hybridisation of cultures" (Rademaker, 2014). Prof. Mutua argues that the right to freedom of religious belief cannot be considered to exist in a "level playing field" in which local cultures can compete with global ones. Rather, the contexts of cultural invasion unfairly privilege global religions, including missionaries making access to education and health services conditional on the "salvation" of "infidels". This echoes arguments by legal scholars that the power and sovereignty dynamics between source and target cultures constitute important factors between "proper" and "improper" proselytism (Stahnke, 1999). The Human Rights Committee has acknowledged that the cultural rights protected under Article 27 depend on the ability of a minority group "to maintain its culture, language or religion".[9] As such, Prof. Mutua argues that the (then Draft) UN Declaration on the Rights of Indigenous People appears to prohibit proselytizing by agents external to the the Indigenous culture in order to create space for Indigenous peoples to maintain their cultures amidst external threats. It follows that one consideration for NLP is whether using Indigenous language translations developed by proselytizing projects constitutes complicity with, and promotion of, projects which might violate international human rights to maintain Indigenous cultures.

5 Recommendations

We take the position that the ACL should strive to be aware of risks of harms to religious communities that NLP research may cause or be complicit with. We suggest that discussions of **Ethical Considerations** should be more common in ACL papers using religious data, and hope the previous section is useful for these. We now discuss how the NLP community might better engage with its cultural standpoints with respect to religious data, and relationships to marginalized religious communities.

Do No Harm. Rogers et al. (2021) provide a useful checklist for responsible data use in NLP, advocating for a principle of "Do No Harm", which includes considering the potential for misuse. Focusing on Indigenous language technology, Schwartz (2022) also argues that the most important concern

[9] General Comment No. 23, UN Doc. ICCPR/C/21/Rev.1/Add.5 (1994), para. 6.2.

is not causing harm. Due to their relationship to deeply held worldviews, we suggest that the use of sacred texts should also prioritize this principle. This entails understanding the harms from the perspectives of the religious communities impacted by the work, for example, a system for automatically translating texts of a global religion into a local Indigenous language should consider possible impacts on those Indigenous communities.

Etic NLP projects. When should an NLP practitioner tread more carefully? We suggest that when handling sacred data from other cultures and religions there will be more morally consequential risks, since the epistemic and axiological uncertainties are greater, if not the aleatoric ones too.

Positionality statements. Linguistic positionality statements have been recommendedn by researchers who are aware of how different priorities and agendas between researchers and language communities can impact projects (e.g., Rolland et al., 2023; Cormier, 2018), and are exemplified in NLP by Ghosh and Caliskan (2023).[10] Similarly, we suggest that religion positionality statements, for NLP research working with religious data, can also provide a useful signal for the NLP community concerning agendas. For example, hypothetically, would a group of non-Muslim researchers have the Muslim community's concerns at heart in developing an automated exegesis system for the Quran?

Transparency. As discussed in Sections 2 and 3, both linguistics and NLP have become dependent to varying degrees on datasets developed by missionary projects. Indeed, even the International Standards Organization's codes for languages, used by many NLP projects, are maintained by SIL (Dobrin and Good, 2009; Morey et al., 2013). We suggest that users of NLP applications should be empowered to make informed decisions that are consistent with their moral worldviews. NLP projects should aim to be transparent about dependencies on NLP resources that have been developed with religious considerations in mind, in line with calls for greater transparency around datasets and models generally (e.g., Bender and Friedman, 2018; Gebru et al., 2021; Mitchell et al., 2019).

Indigenous cultures. Language, religion and culture are not orthogonal or separable. Given the NLP community's skew towards the West (Rungta et al., 2022), the values of, and possible harms to, local communities of diverse cultures are not known by most NLP researchers. For Indigenous NLP projects using translations of sacred texts, we echo calls for more consideration of local colonial contexts, to consider community opinions, and for research to prioritize the needs of Indigenous communities (Bird, 2020; Schwartz, 2022; Alvarado Garcia et al., 2021). Mager et al. (2023) demonstrate one way in which community opinions can be sought regarding NLP projects, and deeper relationships with communities will provide more insights. NLP researchers working with Indigenous languages should become familiar with Indigenous perspectives summarized in Section 4, and with codes of conduct such as the ones published by the Endangered Languages Project.[11] Janke (2021) provides useful guidelines concerning Indigenous Cultural and Intellectual Property (ICIP). With the imminent extinction of many Indigenous languages and religions, we also suggest there may be a role for NLP to play in documenting and maintaining "first-person accounts of what people once believed in and how they talked to and about their gods" (Harrison, 2007, p, 153).

6 Conclusion

This position paper presented a study of the use of religious texts in NLP research, finding that common scenarios are dataset creation and machine translation. We have argued that responsible secularism requires the ACL to engage with concerns about how such NLP activities might impact religious communities, especially the most powerless ones, or might be complicit with projects which do. We provided a detailed account of some of the considerations, with a focus on Indigenous cultures and communities, and suggested how the field can more responsibly engage with questions of religious positionality and cultural standpoints.

7 Researcher Positionality

I live and work in a secular, English-speaking, colonized country of the Global North. I grappled while writing this paper with my lack of first-hand experiential understanding of religion, and thus too with my personal role in arguing for more consideration of global and local religious communities.

[10] As of September 2023, the ACL Anthology has no mentions of "language positionality" nor "linguistic positionality".

[11] https://fpcc.ca/wp-content/uploads/2023/02/CodeOfConduct_Web.pdf

References

David I Adelani, Dana Ruiter, Jesujoba O Alabi, Damilola Adebonojo, Adesina Ayeni, Mofe Adeyemi, Ayodele Awokoya, and Cristina España-Bonet. 2021. The effect of domain and diacritics in Yorùbá-English neural machine translation. *Proceedings of Machine Translation Summit XVIII: Research Track*.

Željko Agić and Ivan Vulić. 2019. JW300: A wide-coverage parallel corpus for low-resource languages. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3204–3210.

Adriana Alvarado Garcia, Juan F Maestre, Manuhuia Barcham, Marilyn Iriarte, Marisol Wong-Villacres, Oscar A Lemus, Palak Dudani, Pedro Reynolds-Cuéllar, Ruotong Wang, and Teresa Cerratto Pargman. 2021. Decolonial pathways: Our manifesto for a decolonizing agenda in HCI research and design. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–9.

Muhammad Huzaifa Bashir, Aqil M Azmi, Haq Nawaz, Wajdi Zaghouani, Mona Diab, Ala Al-Fuqaha, and Junaid Qadir. 2023. Arabic natural language processing for Qur'anic research: A systematic review. *Artificial Intelligence Review*, 56(7):6801–6854.

Emily M Bender and Batya Friedman. 2018. Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics*, 6:587–604.

Steven Bird. 2020. Decolonising speech and language technology. In *Proceedings of the 28th international conference on computational linguistics*, pages 3504–3519.

Steven Bird, Robert Dale, Bonnie J Dorr, Bryan R Gibson, Mark Thomas Joseph, Min-Yen Kan, Dongwon Lee, Brett Powley, Dragomir R Radev, Yee Fan Tan, et al. 2008. The ACL anthology reference corpus: A reference dataset for bibliographic research in computational linguistics. In *LREC*.

Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of "bias" in NLP. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.

Stephanie Russo Carroll, Ibrahim Garba, Oscar L Figueroa-Rodríguez, Jarita Holbrook, Raymond Lovett, Simeon Materechera, Mark Parsons, Kay Raseroka, Desi Rodriguez-Lonebear, Robyn Rowe, et al. 2020. The CARE principles for Indigenous data governance. *Data Science Journal*, 19:43–43.

Rohitash Chandra and Mukul Ranjan. 2022. Artificial intelligence for topic modelling in Hindu philosophy: Mapping themes between the upanishads and the bhagavad gita. *Plos one*, 17(9):e0273476.

Peter A Chew, Steve J Verzi, Travis L Bauer, and Jonathan T McClain. 2006. Evaluation of the Bible as a resource for cross-language information retrieval. In *Proceedings of the workshop on multilingual language resources and interoperability*, pages 68–74.

Gail Cormier. 2018. The language variable in educational research: An exploration of researcher positionality, translation, and interpretation. *International Journal of Research & Method in Education*, 41(3):328–341.

Lise M. Dobrin and Jeff Good. 2009. Practical language development: Whose mission? *Language*, 85(3):619–629.

Patience Epps and Herb Ladley. 2009. Syntax, souls, or speakers? on SIL and community language development. *Language*, 85(3):640–646.

Luciano Floridi and Josh Cowls. 2022. A unified framework of five principles for AI in society. *Machine learning and the city: Applications in architecture and urban design*, pages 535–545.

Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92.

Sourojit Ghosh and Aylin Caliskan. 2023. ChatGPT perpetuates gender bias in machine translation and ignores non-gendered pronouns: Findings across Bengali and five other low-resource languages. *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*.

Courtney Handman. 2009. Language ideologies, endangered-language linguistics, and Christianization. *Language*, 85(3):635–639.

K David Harrison. 2007. *When languages die: The extinction of the world's languages and the erosion of human knowledge*. Oxford University Press.

Dirk Hovy, Shannon L Spruit, Margaret Mitchell, Emily M Bender, Michael Strube, and Hanna Wallach. 2017. Proceedings of the first ACL workshop on ethics in natural language processing. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*.

Ben Hutchinson, Negar Rostamzadeh, Christina Greer, Katherine Heller, and Vinodkumar Prabhakaran. 2022. Evaluation gaps in machine learning practice. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1859–1876.

Ayyoob ImaniGooghari, Peiqin Lin, Amir Hossein Kargaran, Silvia Severini, Masoud Jalili Sabet, Nora Kassner, Chunlan Ma, Helmut Schmid, André FT Martins, François Yvon, et al. 2023. Glot500: Scaling multilingual corpora and language models to 500

languages. *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Terri Janke. 2021. *True tracks: Indigenous cultural and intellectual property principles for putting self-determination into practice*. University of New South Wales Press.

Sneha Kudugunta, Isaac Caswell, Biao Zhang, Xavier Garcia, Christopher A Choquette-Choo, Katherine Lee, Derrick Xin, Aditya Kusupati, Romi Stella, Ankur Bapna, et al. 2023. MADLAD-400: A multilingual and document-level large audited dataset. *arXiv preprint arXiv:2309.04662*.

Manuel Mager, Elisabeth Mager, Katharina Kann, and Ngoc Thang Vu. 2023. Ethical considerations for machine translation of Indigenous languages: Giving a voice to the speakers. *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics Volume 1: Long Papers*.

Stephen Mayhew, Chen-Tse Tsai, and Dan Roth. 2017. Cheap translation for cross-lingual named entity recognition. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, pages 2536–2545.

Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 220–229.

Stephen Morey, Mark W Post, and Victor A Friedman. 2013. The language codes of ISO 639: A premature, ultimately unobtainable, and possibly damaging standardization.

Till Mostowlansky and Andrea Rota. 2020. Emic and etic. In Felix Stein, editor, *The Open Encyclopedia of Anthropology*. University of Cambridge.

Makau Mutua. 2004. Proselytism and cultural integrity. In *Facilitating Freedom of Religion or Belief: A Deskbook*, pages 651–668. Brill Nijhoff.

National Health and Medical Research Council (Australia). 2018. *Ethical conduct in research with Aboriginal and Torres Strait Islander Peoples and communities: Guidelines for researchers and stakeholders*. National Health and Medical Research Council.

Robert Cummings Neville. 2018. *Defining religion: Essays in philosophy of religion*. State University of New York Press.

Shrimai Prabhumoye, Brendon Boldt, Ruslan Salakhutdinov, and Alan W Black. 2020. Case study: Deontological ethics in NLP. *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.

Vineel Pratap, Andros Tjandra, Bowen Shi, Paden Tomasello, Arun Babu, Sayani Kundu, Ali Elkahky, Zhaoheng Ni, Apoorv Vyas, Maryam Fazel-Zarandi, et al. 2023. Scaling speech technology to 1,000+ languages. *arXiv preprint arXiv:2305.13516*.

Mahima Pushkarna, Andrew Zaldivar, and Oddur Kjartansson. 2022. Data cards: Purposeful and transparent dataset documentation for responsible AI. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1776–1826.

Laura Rademaker. 2014. Language and Australian Aboriginal history; Anindilyakwa and English on Groote Eylandt. *History Australia*, 11(2):222–239.

John Rawls. 1971. *A theory of justice*. Cambridge, MA: Harvard University Press.

Philip Resnik, Mari Broman Olsen, and Mona Diab. 1999. The Bible as a parallel corpus: Annotating the 'book of 2000 tongues'. *Computers and the Humanities*, 33:129–153.

Anna Rogers, Tim Baldwin, and Kobi Leins. 2021. Just what do you think you're doing, Dave? A checklist for responsible data use in NLP. *Findings of the Association for Computational Linguistics: EMNLP 2021*.

Louise Rolland, Hannah M King, and Pernelle Lorette. 2023. Methodological implications of participant and researcher multilingualism: making language dynamics visible. *Journal of Multilingual and Multicultural Development*, 44(8):645–656.

Mukund Rungta, Janvijay Singh, Saif M Mohammad, and Diyi Yang. 2022. Geographic citation gaps in NLP research. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.

Lane Schwartz. 2022. *Primum Non Nocere*: Before working with Indigenous data, the ACL must confront ongoing colonialism. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, volume 2.

Melford E Spiro. 2013. Religion: Problems of definition and explanation. In *Anthropological approaches to the study of religion*, pages 85–126. Routledge.

Tad Stahnke. 1999. Proselytism and the freedom to change religion in international human rights law. *BYU L. Rev.*, page 251.

Karaitiana Taiuru. 2021. Kaitiakitanga Māori data sovereignty licences. Retrieved from https://www.taiuru.maori.nz/maori-data-sovereignty-licences/.

Langdon Winner. 2017. Do artifacts have politics? In *Computer ethics*, pages 177–192. Routledge.