

Toward an Atlas of Cultural Commonsense for Machine Reasoning

Anurag Acharya,¹ Kartik Talamadupula,² Mark A Finlayson¹

¹ School of Computing and Information Sciences
Florida International University

² IBM Research

{aacharya, markaf}@fiu.edu, krtalamad@us.ibm.com

Abstract

Existing commonsense reasoning datasets for AI and NLP tasks fail to address an important aspect of human life: cultural differences. We introduce an approach that extends prior work on crowdsourcing commonsense knowledge by incorporating differences in knowledge that are attributable to cultural or national groups. We demonstrate the technique by collecting commonsense knowledge that surrounds six fairly universal rituals—birth, coming-of-age, marriage, funerals, new year, and birthdays—across two national groups: the United States and India. Our study expands the different types of relationships identified by existing work in the field of commonsense reasoning for commonplace events, and uses these new types to gather information that distinguish the identity of the groups providing the knowledge. It also moves us a step closer towards building a machine that doesn't assume a rigid framework of universal (and likely Western-biased) commonsense knowledge, but rather has the ability to reason in a contextually and culturally sensitive way. Our hope is that cultural knowledge of this sort will lead to more human-like performance in NLP tasks such as question answering (QA) and text understanding and generation.

1 Introduction

In the past few years, there have been major advancements in the field of question answering (QA) systems (Gan and Ng 2019; Fan and Ferrucci 2019; Qu et al. 2019; Zafar et al. 2020), in which researchers have looked at different ways in which these systems can be made more accurate and human-like in both their responses as well as their methodology. Incorporating commonsense knowledge and reasoning into NLP systems is one such area of recent focus (Tandon et al. 2018; Tandon, Varde, and de Melo 2018; Trinh and Le 2018; Merkhofer et al. 2018), and a large body of recent work has focused on the creation, curation, and use of large-scale commonsense knowledge bases and knowledge graphs (Sap et al. 2019b; Bosselut et al. 2019). Importantly, these types of knowledge acquisition efforts have a long history in and have been of great use to a wide variety of AI systems (Shi et al. 2017; Olteanu, Varol, and Kiciman 2017; Sap et al. 2019a; Liu et al. 2020).

The importance of commonsense knowledge bases and repositories is clear from the volume of recent work that

makes use of resources such as ConceptNet (Speer, Robyn 2020; Speer, Chin, and Havasi 2016; Speer and Havasi 2012) to imbue NLP systems with worldly knowledge obtained from humans. A key recent contribution along these lines was ATOMIC (Sap et al. 2019b), which tackles the task of incorporating commonsense reasoning into NLP tasks by generating an atlas of “if-then” rules that taken together produce behavior akin to commonsense reasoning. Work such as ATOMIC and COMET (Bosselut et al. 2019) has made commonsense knowledge more accessible to the current generation of NLP systems; the progress and pitfalls of this work have been cataloged recently (Sap et al. 2020).

One glaring omission in all of this prior work has been the lack of focus on context-contingent aspects of commonsense knowledge; that is, most prior work views commonsense as a universal monolith. While some events included in prior work are not variable across groups—like reading a book or breaking a window, for instance—many events are variable, and here we focus on one highly relevant type of context-specific commonsense knowledge, namely cultural commonsense. Consisting of ritualistic, geographical, and social knowledge, cultural commonsense plays a large but hidden role in humans' day-to-day social interactions.

For example, let us consider a very simple social setting: *You are invited to a wedding. How long do you expect to be gone for, and how many people do you think will be there?* For most people in the United States or the wider Western world, the answer would probably be a few hours; probably half a day, starting in the early afternoon; and somewhere around a 100 people. However, for many people in India, the obvious answer is that you will probably have to lay aside several days for the whole event, and anywhere between several hundred to over a thousand people will attend. Such socially-conditioned knowledge is inherently obvious to people from the respective cultures, and hints at the differences in commonsense knowledge across cultural and social settings, particularly when it comes to ritualistic practices.

We build upon prior work on systematizing commonsense knowledge for use in NLP tasks by demonstrating a proof-of-concept scheme for gathering *cultural commonsense* in a format similar to previous approaches like ATOMIC. Specifically, we start by surveying the extensive prior literature on cultural knowledge and ritual practices, and select a short list of six rituals to focus on for our study. We select two differ-

ent national groups that are diverse in their ritualistic practices, and conduct a pilot experiment via a survey. We report on the results of the survey, and showcase what a truly cultural commonsense knowledge repository might look like. We hope that this work spurs future research on incorporating cultural and social commonsense knowledge into NLP systems across a wide range of tasks.

2 Related Work

While the concept of incorporating cultural knowledge into commonsense is novel, there have been several previous attempts at laying the groundwork for it. We build on the ATOMIC (Sap et al. 2019b) system and knowledge repository, where crowdsourced commonsense information was used to build an atlas for if-then reasoning. ATOMIC builds a knowledge graph containing inferential knowledge regarding 24,000 “short events.” The dataset is then used on the social question answering system SocialIQA (Sap et al. 2019c), which shows an increase in performance using the commonsense knowledge from ATOMIC.

Another prominent effort is the AI2 Reasoning Challenge (ARC) (Clark et al. 2018). ARC consisted of a dataset of almost 8,000 science questions in English. This dataset was split into the *Easy* set and the *Challenge* set; the *Challenge* set consisted of questions that neither a retrieval-based algorithm nor a word co-occurrence algorithm were able to answer correctly. Later, Boratko et al. (2018) more precisely analyzed the ARC knowledge, defining seven knowledge types and nine reasoning types, as well as triple annotating 192 ARC questions.

In addition to ATOMIC and ARC, there have been several other efforts in the field of commonsense reasoning. Rashkin et al. (2018) built a system called Event2mind—a commonsense inference system on events, intents, and reactions. Furthermore, Speer, Chin, and Havasi (2016) have built an updated version of ConceptNet (Liu and Singh 2004), a practical commonsense reasoning toolkit, which is now a multilingual graph of general knowledge. Another relevant work in the field is Webchild 2.0, a fine-grained commonsense knowledge distillation (Tandon, De Melo, and Weikum 2017). Several of the other important pieces of work done in the field of commonsense have been reviewed by Davis and Marcus (2015), which lays out the uses, successes, challenges, approaches and possible future work in the field of commonsense reasoning. Apart from these applications, Gordon and Hobbs (2017) lays out a formal theory of commonsense psychology and how people assume others think, while Lake et al. (2017) put forward their argument as to how we can go about building machines that learn and think like people.

There have also been several efforts to facilitate Question Answering (QA) systems and commonsense knowledge by building datasets (Talmor et al. 2018; Dunn et al. 2017; Joshi et al. 2017). An important one is the Stanford Question Answering Dataset (SQuAD) (Rajpurkar et al. 2016), consisting of 100,000+ questions and a reading comprehension dataset. They contrast three types of tasks: reading comprehension (RC; read a passage, select a span that answers); Open-domain QA (answer a question from a large set of

documents); and Cloze datasets (predict a missing word in a passage). Similarly Sap et al. (2019c) introduces a benchmark dataset for social and emotional commonsense reasoning and the MCTest dataset (Richardson, Burges, and Renshaw 2013), which consists of multiple-choice reading comprehension questions comprised of short (150–300 words) fictional stories which were targeted at seven-year olds.

3 Cultural Knowledge and Rituals

The primary focus of this work is to encode cultural knowledge as an essential part of commonsense knowledge. To do this, we need to first find ways to collect cultural knowledge across a variety of groups. This is not straightforward, since cultural knowledge varies dramatically. Prior literature from the study of cultural groups suggests that it is hard enough just to define “culture” given the many complexities and nuances of group differences; it is even harder to identify the knowledge that comes along with it and to differentiate across cultural groups. Furthermore, as cultural norms and practices are so varied across various groups, it is hard to find knowledge structures that can be clearly compared across different groups. The first task, therefore, is to find topics that are relatively common across different groups that can be compared in a like-for-like fashion.

3.1 Rituals

We first detail our search for cultural markers that can be consistently compared across different groups. After a survey of the extant literature, we decided to focus on *rituals* that are commonly found across cultures. There are several definitions of *ritual*, and the concept often comes entangled with religion and rites (Braun and McCutcheon 2000, p. 259–262). For the purposes of this work, we define ritual as a “culturally defined set of behavior” (Leach 1968).

While a detailed review of the theory of rituals is beyond the scope of this paper, we can list several reasons for focusing on rituals as indicators of cultures. First, rituals are well-studied; Scholars have identified not just specific rituals, but also the genres and types of rituals and their variations across cultures, etc. (Durkheim 1915; Turner 1973; Bell 1992, 1997). There is also a rich body of literature on the analysis of cultural practices for various rituals across several cultures, clearly indicating a strong concomitance between rituals and cultures (Gray 1979; Smith 1986; Underhill 2000).

While there great variety in the conceptualization of ritual, the most important taxonomies of ritual show basic agreement on their core categories. Rituals can be classified in several ways: one of the more widely accepted categorizations of rituals by Bell (1992, 1997) presents a compromise of six categories that are not necessarily mutually exclusive. These categories are: (i) *rites of passage*, a.k.a. *life-cycle rites*; (ii) *calendrical and commemorative rites*; (iii) *rites of exchange and communion*; (iv) *rites of affliction*; (v) *rites of feasting, fasting, and festivals*; and (vi) *political rituals*.

3.2 Selecting Rituals

Given the vast numbers of rituals that can fall into the six categories previously outlined, and the variance in the ex-

tents of their observance across cultures, another crucial decision is in selecting specific rituals as markers of cultural knowledge. There were several factors that we had to consider while making this decision. First and foremost, we needed activities whose identifying names are used in multiple cultures. For example, a ritual like *Passover* cannot be used since it is highly specific to Judaism and Jewish cultures. It would make no sense to ask “*What are your cultural practices with regards to Passover?*” of a Hindu or a Muslim. Furthermore, the more specific the names, the easier it is for a QA system to associate its knowledge to that activity. As an example, “*Catholic Christmas Mass*” is a highly denomination- and group-specific ritual, and will exhibit very little variation across cultures. Second, in order to ease the data collection process, the selected rituals needed to have fairly concise and telegraphic names. For instance, it is confusing to probe a participant in a study about “*the kinds of things you do before a sports game popular in your culture*”; even though this is an activity that is fairly widespread and at the same time variable across cultures. Instead, we are seeking rituals that can be described in just a few words and bring a very specific activity or event to mind. Our choice of rituals is intended to primarily ease the collection of crowd-sourced data—we thus pick activities that may have different practices across different cultural groups but are likely to be found in all of them. We take guidance from the analysis of Bell (1997) and use the following six rituals as our target rituals in this work:

1. **Wedding:** wed, marriage, marry, matrimony, nuptials, wedlock, union, hymeneals (*rite of passage*)
2. **Funeral:** funerary, burial, cremation, interment, entombment, obsequy (*rite of passage*)
3. **Coming of Age:** becoming a man, becoming a woman, manhood, womanhood, adulthood (*rite of passage*)
4. **Birth:** childbirth, delivery, birthing, childbearing, parturition, nativity (*rite of passage*)
5. **New Year** (*calendrical rite*)
6. **Birthday:** name day, natal day (*rite of passage*)

These rituals all have the advantage that across cultural groups there are limited number of ways of naming or expressing them, and the meaning is evident to most subjects answering our survey.

4 Methodology

In this section, we outline the process of gathering the ritual-based cultural knowledge. We first describe the target cultures, followed by the details of our pilot experiment; we explain our data collection method on Amazon Mechanical Turk (MTurk); finally we describe the survey questionnaire that we used. We share these details both to explain our method as well as to enable reproducibility of the work.

4.1 Selecting Target Cultures

For this study, we focused on two specific target groups: Americans (people from the United States of America) and Indians (people from India). We chose these two groups for

a variety of reasons. First, both the countries use English as one of their major languages, either officially or unofficially, and since our study was to be conducted in English, this was a key requirement. Second, Amazon Mechanical Turk has a high presence of workers from both these countries (Difallah, Filatova, and Ipeirotis 2018). Since the bulk of the data collection was to be done via Amazon’s Mechanical Turk platform, it was essential that we considered the demographics of the crowd-workers. Moreover, these groups allowed us to set up a unique contrast and high degree of cultural variation between the two groups. Apart from providing contrasts against each other, the fact that the United States and India are large and diverse countries consisting of various cultures allows us to capture a varied amount of data within the groups themselves.

4.2 Data Collection: Pilot Experiment

For the pilot experiment, we collected data from a small number of participants for three rituals: coming of age, wedding, and death rites/funeral. We collected two unique sets of responses per ritual per culture, for a total of 12 unique responses. The identification of cultural group membership was done via self-identification by the participants, based on a demographic questionnaire that preceded the main survey. The participants took the survey using an online form. The actual survey consisted of a series of questions that were modifications of the ATOMIC (Sap et al. 2019b) question set (see Section 4.4 for details). The questions remained constant across rituals and cultures with only the initial prompt changing to keep the method as consistent as possible. The survey was conducted asynchronously and participants were compensated for their time.

4.3 Data Collection: Amazon MTurk

Once our small pilot study affirmed our idea, we moved ahead with collecting data on a slightly larger scale. We used Amazon Mechanical Turk (MTurk) to collect our second round of data. We set up separate tasks on MTurk for United States (US) participants and Indian (IN) participants, and geo-restricted the tasks to workers from the respective countries. We also restricted the tasks to *Master* workers—workers with a work approval rate of 90% or more. We set up a system of auto-generated survey codes that linked responses to MTurk workers without having to collect any Personally Identifiable Information (PII); these codes were used to filter out spam entries. The workers that provided spam responses were banned from the tasks. Overall, we collected a total of 32 useful data points for Indian participants, and 33 for US participants; these were divided roughly equally across rituals, with at least 5 responses per ritual.

4.4 Survey Questionnaire

The survey questionnaire was divided into two sections: the self-identification questionnaire, and the survey proper.

Self-identification Questions We based the self-identification questionnaire primarily on the *College Students Knowledge and Belief* (Clark and Barrows 1981) questionnaire. Portions of the questionnaire relevant to our

Answer the following questions about a typical wedding/marriage in your society

Where does this event typically happen?

When does this event typically happen?

How long does this event typically last?
 --Choose Unit-- ▾

How many people typically participate in an event like this?

Who are the important people involved in this event?

Person 1

Person 2

Person 3

Is one or more of the important people the focus of this event?

Yes No

Select the people who are the focus of this event.

Person 1
 Person 2
 Person 3

Figure 1: Layout of the main survey form as seen by the survey participants.

current study were modified to suit our purposes, and the wording changed to be consistent with contemporary terminology. This self-identification questionnaire was further modified to prioritize the targeted countries with regard to language and religion. The options for these two fields were based on the most likely answers given the demographics of those populations. The updated questionnaire asked participants about their place of birth, languages used, and religion; it also asked these questions about their primary caregivers in order to ascertain the degree to which the participants were immersed into their identified culture. While for this work we have considered only nationality as the marker for culture, moving forward we would want to use a broader set of factors such as these to define the group.

Event-specific Questions For the main survey, a prompt specifying the ritual event under consideration was first shown, and the participants were asked a set of questions pertaining to the specific event; see the questions asked and the layout of the form in Figure 1 for details.

Person-specific Questions After the event-specific questions, the participants were asked to answer questions that involved the specific persons mentioned in the events (as part of their responses). This part of the survey was adapted from ATOMIC (Sap et al. 2019b). Questions were divided into three temporal categories: before the event, during the event, and after the event. We considered four types of questions: (1) intent & reaction; (2) need & want; (3) effects; and (4) attributes. These four types evolved into 11 ques-

BEFORE THE EVENT
Does this person typically have an intent in causing the event? What is this person’s typical intent in causing the event? Does this person typically need to do anything before this event? What does this person typically need to do before this event?
DURING THE EVENT
Does this person typically use something during the event? What things does this person typically use during this event?
AFTER THE EVENT
How would this person be described as a consequence of the event? Does this person typically want to do something after this event? What does this person typically do after this event? What is the typical effect of the event on this person? What does this person typically feel after this event?

Table 1: Questions asked of the survey participants for each *important* person associated to an event in their responses.

tions on the survey form, as shown in Figure 1. In terms of the presentation of the questions, the terms `PersonX` or `PersonY` were replaced by the actual names that participants provided (in the Person fields) in order to make the questions feel more natural to the survey participants. These person-specific questions were repeated for each person that the survey participant deemed “important” to a given event.

5 Results

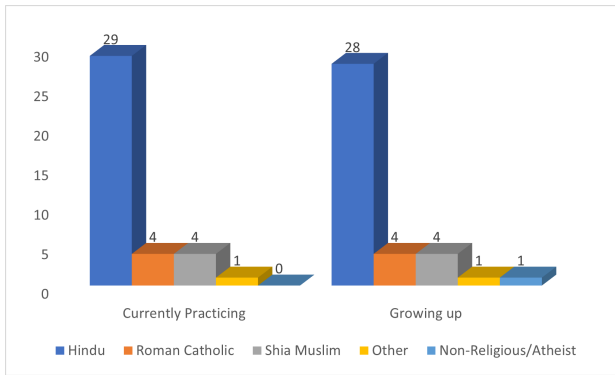
In this section, we detail the findings of our study. We first report on the demographics of the survey participants, particularly with an eye towards cultural background, and then recount and discuss the responses to the survey.

5.1 Demographics

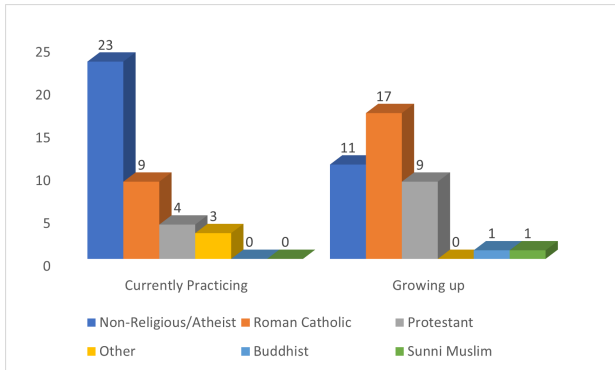
We obtained a total of 77 unique survey responses with each participant limited to one response per ritual per culture. Only information that was deemed necessary for the purposes of the experiment was collected from participants, with an aim to avoid collecting PII as far as possible.

Geography Out of the 77 responses, 38 were from India and 39 from the US. While all the participants that identified as American (from the US) resided in the US; in the case of Indian participants, three people said they currently resided in the US.

Religion Since many rituals have a basis in religion—to the extent that they are often intertwined (Goody 1961; Geertz and Banton 1966; Bell 1992)—it is important to ensure a diversity of religious practice among the respondents. Of the participants in the study, 29 said they practiced Hinduism; 17 said they practiced some form of Christianity; four Islam; four other religions; and 23 said they did not practice any religion and/or were atheist ($N = 77$). Among



(a) Religion Data for the participants from India.



(b) Religion Data for the participants from the US.

Figure 2: Variation in the religion practiced by participants growing up versus currently for both cultures.

participants who identified as Indian, there was significantly less variation in religion compared to the participants who identified as Americans. There was also significantly more change in religion practiced over a lifetime for US participants as compared to IN. This is illustrated in Figure 2.

Language Among the $N = 77$ participants, 39 identified their native language as English; 17 Tamil; seven Hindi; five Telugu; four Urdu; and one each as five other languages.

5.2 Qualitative Analysis of Survey Responses

Over both the pilot and MTurk studies, we collected a total of 77 responses, with a minimum of 5 responses per ritual (for six rituals, see Section 3.2 for a list) per cultural group. Since there are not enough data points yet in our pilot experiment to conduct a quantitatively significant study, we instead present a qualitative analysis of the data collected thus far.

The results show us that while rituals have some common features across cultures, they can also have significant variations to the point where the *common knowledge* would be noticeably different. Let us look at one significant difference seen in the responses: for the *wedding* ritual, participants from the US said the bride would focus on the wedding planning part of the event, like dresses and so forth; while the Indian participants focused on the cultural aspects of the wedding, as well as the fact that the bride might have

to get to know the groom’s family, and possibly the groom himself, as illustrated in Figure 6. It would be extremely unlikely for a bride to not know the groom’s family, let alone the groom himself, in a typical US wedding; while this is still reasonably prevalent in Indian society. This is an excellent example of the type of knowledge that is collected by our work, where a machine can now leverage this information as commonsense knowledge that is culturally sensitive and correct. This also suggests that with more data, more such variations in the way rituals are conducted across cultures can be documented and understood by NLP systems.

While the data showed that the person-roles involved in rituals were more or less similar across cultures, the order of importance indicated was noteworthy, and shown in Table 2. For birth rituals, participants in the US considered the birth itself the main event, and said the Doctor (physician) was an important person (after the parents); while Indian participants focused more on the family. It is also worth noting that US participants put Mother before Father, while Indian participants reversed the priority order. All participants gave similar responses in terms of the important person-roles for a wedding: the bride, the groom, and their parents. An important difference was that Americans identified the Bride’s mother as the more important person, while Indians picked the Groom’s mother for the same. Likewise, for a funeral, more priority was given to the pastor or priest by US participants; unlike Indian participants, who considered the family to be the more important participants.

Another set of major differences were seen in terms of the durations of the rituals. While the duration of the other four rituals were more or less comparable, weddings and funerals had significant differences across cultures. For the *wedding* ritual, participants from the US said that the ritual typically lasted a few hours, while Indian participants responded by saying that weddings lasted multiple days. The difference was even more striking for funerals, which US participants said lasted only a few hours; while Indian participants reported funeral rites lasted 13 days, as shown in Figure 3(a).

One other key difference observed was the number of people that participated in each ritual. With the exception of *coming-of-age* rituals, for which the number of participants were reported to be similar, the rituals in India seemed to involve a lot more people than in the US. While the difference in numbers was already pronounced for other rituals (Indian funerals seem to have *twice* as many people as US ones), the striking difference is weddings, where Indian ones had several hundred guests, while US ones remained below 100, as shown in Figure 3(b).

These findings validate our expectation that rituals can give us a peek into cultures and how they vary, and that commonsense knowledge cannot truly be complete without including cultural nuances.

6 A Cultural Commonsense Graph

The main motivation behind this study is to eventually distill the knowledge thus gathered into a form useful to current NLP systems. Specifically, our hope is to transform the cultural knowledge thus collected into a resource that can be used to produce more human-like performance in NLP

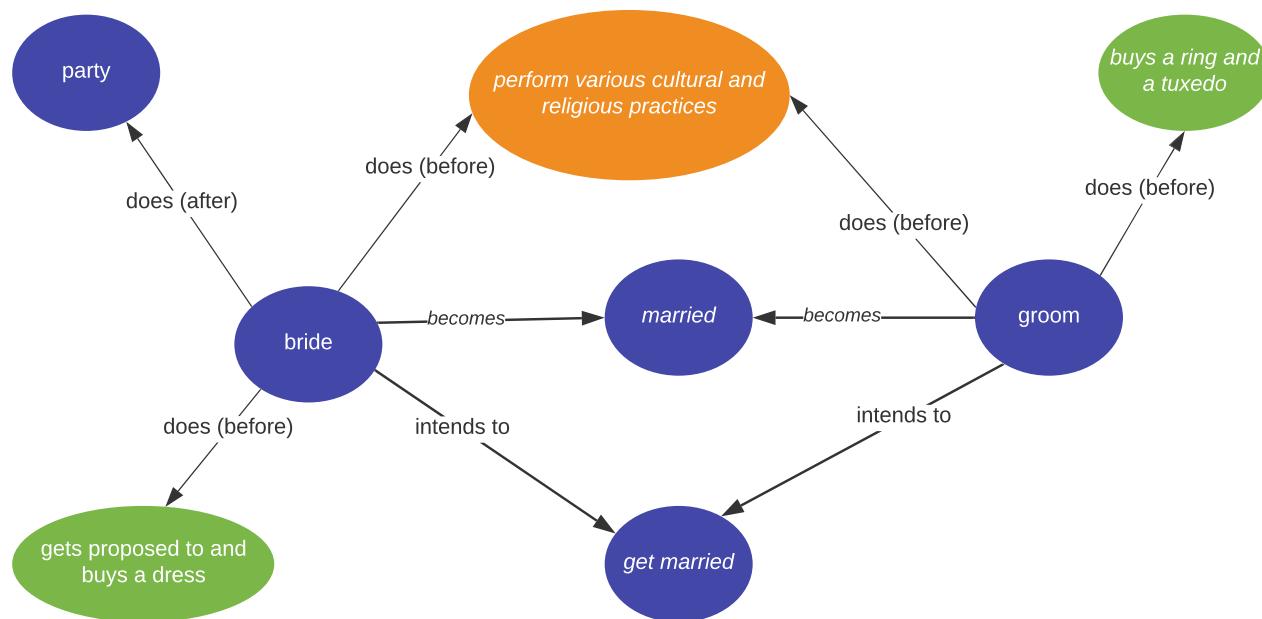


Figure 5: A graph for the subset of the data collected for the *wedding* ritual. The nodes in orange show responses of Indian participants; those in green show responses of US participants; nodes in purple are common to both. The thickness of the edge indicates the frequency of the relationship. The figure has been manually edited for grammar, and is intended for illustration.

to exhibit human-like performance by addressing gaps in their current knowledge. We showcased some interesting qualitative results. The task of injecting cultural sensitivity into commonsense reasoning, while being crucial to developing a true human-like AI, has not been previously explored in the field. We exposed this gap and in order to bridge it performed the difficult task of choosing suitable cultural markers that would work within existing frameworks of commonsense knowledge. Moving forward, it is important that we carefully assess how culturally variable different events are, as well as measure the significance of any differences for these events across culture as we look to scale up this work. Another direction for future work is to explore the various knowledge representation techniques to find the best way to represent what will likely be multi-dimensional data. While the data we have collected thus far is too small to be used to directly improve the performance of QA or other NLP tasks, the approach here allows for scale-up into a full dataset. We are currently in that process, and we are also evaluating existing datasets to design benchmarks that require the use of cultural commonsense.

8 Acknowledgements

This project was supported in part by an 2019 IBM Career Award to Prof. Finlayson, as well as in part by DARPA contract FA8650-19-C-6017.

References

Bell, C. 1992. *Ritual theory, ritual practice*. Oxford: Oxford University Press.

Bell, C. M. 1997. *Ritual: Perspectives and dimensions*. Oxford: Oxford University.

Boratko, M.; Padigela, H.; Mikkilineni, D.; Yuvraj, P.; Das, R.; McCallum, A.; Chang, M.; Fokoue-Nkoutche, A.; Kapanipathi, P.; Mattei, N.; Musa, R.; Talamadupula, K.; and Witbrock, M. 2018. A Systematic Classification of Knowledge, Reasoning, and Context within the ARC Dataset. In *Proceedings of the Workshop on Machine Reading for Question Answering (MRQA) at ACL 2018*, 60–70. Melbourne, Australia.

Bosselut, A.; Rashkin, H.; Sap, M.; Malaviya, C.; Celikyilmaz, A.; and Choi, Y. 2019. COMET: Commonsense Transformers for Automatic Knowledge Graph Construction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4762–4779. Florence, Italy.

Braun, W.; and McCutcheon, R. 2000. *Guide to the Study of Religion*. London: Bloomsbury.

Clark, J.; and Barrows, T. 1981. College students’ knowledge and beliefs: A survey of global understanding. Technical report, Educational Testing Service.

Clark, P.; Cowhey, I.; Etzioni, O.; Khot, T.; Sabharwal, A.; Schoenick, C.; and Tafjord, O. 2018. Think you have solved question answering? try ARC, the AI2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.

Davis, E.; and Marcus, G. 2015. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Communications of the ACM* 58(9): 92–103.

Difallah, D.; Filatova, E.; and Ipeirotis, P. 2018. Demographics and dynamics of mechanical Turk workers. In *Proceedings of the 11th ACM International Conference on Web Search and Data Mining*, 135–143. Los Angeles, CA.

- Dunn, M.; Sagun, L.; Higgins, M.; Guney, V. U.; Cirik, V.; and Cho, K. 2017. Searchqa: A new q&a dataset augmented with context from a search engine. *arXiv preprint arXiv:1704.05179* .
- Durkheim, E. 1915. *The elementary forms of the religious life*. London: G. Allen & Unwin.
- Ellson, J.; Gansner, E. R.; Koutsofios, E.; North, S. C.; and Woodhull, G. 2004. Graphviz and dynagraph—static and dynamic graph drawing tools. In *Graph drawing software*, 127–148. Berlin: Springer.
- Fan, J. J.; and Ferrucci, D. A. 2019. Scoring candidates using structural information in semi-structured documents for question answering systems. US Patent 10,223,441.
- Gan, W. C.; and Ng, H. T. 2019. Improving the robustness of question answering systems to question paraphrasing. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 6065–6075. Florence, Italy.
- Geertz, C.; and Banton, M. 1966. *Religion as a cultural system*. London: Tavistock.
- Goody, J. 1961. Religion and ritual: the definitional problem. *The British Journal of Sociology* 12(2): 142–164.
- Gordon, A. S.; and Hobbs, J. R. 2017. *A formal theory of commonsense psychology: How people think people think*. Cambridge: Cambridge University Press.
- Gray, J. N. 1979. Keep the Hom fires burning: sacrifice in Nepal. *Social Analysis: The International Journal of Social and Cultural Practice* 1(1): 81–107.
- Joshi, M.; Choi, E.; Weld, D. S.; and Zettlemoyer, L. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551* .
- Lake, B. M.; Ullman, T. D.; Tenenbaum, J. B.; and Gershman, S. J. 2017. Building machines that learn and think like people. *Behavioral and Brain Sciences* 40: e253.
- Leach, E. R. 1968. *A runaway world?* Oxford University Press.
- Liu, H.; and Singh, P. 2004. ConceptNet—a practical commonsense reasoning tool-kit. *BT technology journal* 22(4): 211–226.
- Liu, W.; Zhou, P.; Zhao, Z.; Wang, Z.; Ju, Q.; Deng, H.; and Wang, P. 2020. K-BERT: Enabling Language Representation with Knowledge Graph. In *Proceedings of the 34th Annual Conference on Artificial Intelligence (AAAI 2020)*, 2901–2908.
- Merkhofer, E.; Henderson, J.; Bloom, D.; Strickhart, L.; and Zarella, G. 2018. Mitre at semeval-2018 task 11: Commonsense reasoning without commonsense knowledge. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, 1078–1082.
- Olteanu, A.; Varol, O.; and Kiciman, E. 2017. Distilling the outcomes of personal experiences: A propensity-scored analysis of social media. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, 370–386. Portland, OR.
- Qu, C.; Yang, L.; Croft, W. B.; Scholer, F.; and Zhang, Y. 2019. Answer interaction in non-factoid question answering systems. In *Proceedings of the 2019 Conference on Human Information Interaction and Retrieval*, 249–253. Glasgow, Scotland.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250* .
- Rashkin, H.; Sap, M.; Allaway, E.; Smith, N. A.; and Choi, Y. 2018. Event2mind: Commonsense inference on events, intents, and reactions. *arXiv preprint arXiv:1805.06939* .
- Richardson, M.; Burges, C. J.; and Renshaw, E. 2013. Mctest: A challenge dataset for the open-domain machine comprehension of text. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, 193–203. Seattle, WA.
- Sap, M.; Gabriel, S.; Qin, L.; Jurafsky, D.; Smith, N. A.; and Choi, Y. 2019a. Social bias frames: Reasoning about social and power implications of language. *arXiv preprint arXiv:1911.03891* .
- Sap, M.; Le Bras, R.; Allaway, E.; Bhagavatula, C.; Lourie, N.; Rashkin, H.; Roof, B.; Smith, N. A.; and Choi, Y. 2019b. ATOMIC: An Atlas of Machine Commonsense for If-then Reasoning. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, 3027–3035. Honolulu, HI.
- Sap, M.; Rashkin, H.; Chen, D.; Le Bras, R.; and Choi, Y. 2019c. Social IQA: Commonsense Reasoning about Social Interactions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4453–4463. Hong Kong, China.
- Sap, M.; Shwartz, V.; Bosselut, A.; Choi, Y.; and Roth, D. 2020. Commonsense Reasoning for Natural Language Processing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, 27–33. Online.
- Shi, F.; Chen, L.; Han, J.; and Childs, P. 2017. A data-driven text mining and semantic network analysis for design information retrieval. *Journal of Mechanical Design* 139(11): 111402.
- Smith, B. K. 1986. The unity of ritual. *Indo-Iranian Journal* 29(2): 79–96.
- Speer, R.; Chin, J.; and Havasi, C. 2016. Conceptnet 5.5: An open multilingual graph of general knowledge. *arXiv preprint arXiv:1612.03975* .
- Speer, R.; and Havasi, C. 2012. Conceptnet 5. *Tiny Transactions of Computer Science* .
- Speer, Robyn. 2020. ConceptNet in Context. Invited talk at the Reasoning for Complex QA (RCQA) Workshop at AAAI 2020.
- Talmor, A.; Herzig, J.; Lourie, N.; and Berant, J. 2018. CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge. *CoRR* abs/1811.00937. URL <http://arxiv.org/abs/1811.00937>.
- Tandon, N.; De Melo, G.; and Weikum, G. 2017. Webchild 2.0: Fine-grained commonsense knowledge distillation. In *Proceedings of ACL 2017, System Demonstrations*, 115–120. Vancouver, Canada.
- Tandon, N.; Mishra, B. D.; Grus, J.; Yih, W.-t.; Bosselut, A.; and Clark, P. 2018. Reasoning about actions and state changes by injecting commonsense knowledge. *arXiv preprint arXiv:1808.10012* .
- Tandon, N.; Varde, A. S.; and de Melo, G. 2018. Commonsense knowledge in machine intelligence. *ACM SIGMOD Record* 46(4): 49–52.
- Trinh, T. H.; and Le, Q. V. 2018. A simple method for commonsense reasoning. *arXiv preprint arXiv:1806.02847* .
- Turner, V. W. 1973. Symbols in African ritual. *Science* 179(4078): 1100–1105.
- Underhill, A. 2000. An analysis of mortuary ritual at the Dawenkou site, Shandong, China. *Journal of East Asian Archaeology* 2(1): 93–127.
- Zafar, H.; Dubey, M.; Lehmann, J.; and Demidova, E. 2020. IQA: Interactive query construction in semantic question answering systems. *Journal of Web Semantics* 64: 100586.

Appendix A: Ethical Considerations

We were extremely cognizant of the ethical aspects of our study. The details of the study were reviewed by the Institutional Review Board (IRB) of Florida International University. The proposed study was deemed *Exempt Research*, and given approval to proceed. No step of the study—excepting internal preparations—were carried out prior to obtaining this approval. No personally identifiable information was collected from the participants, and all IRB requirements were met and followed during the course of the study.

Another issue is whether a study such as this might exacerbate existing stereotypes. We agree that this is a concern, and requires careful consideration in future work.

Appendix B: Additional Figures and Tables

<i>US</i>	<i>India</i>	<i>US</i>	<i>India</i>
Birth		Birthday	
Mother	Father	Self	Parents
Father	Mother	Partner	Family
Doctor	Family	Family	Relatives
New Year		Coming-of-age	
Spouse	Friends	Self	Parents
Parents	Spouse	Parents	Family
Friends	Family	Siblings	Relatives
Wedding		Funeral	
Bride	Bride	Spouse	Father
Groom	Groom	Pastor/Priest	Son
Bride's Mother	Groom's Mother	Parents	Family

Table 2: Important people for each ritual by culture.

Before a wedding,
the bride...



... plans the wedding



... gets to know groom's family

Figure 6: The differences in expectations of a bride in the US versus in India.