


Please cite the Published Version

Mohamed, Emad and Sarwar, Raheem  (2024) Exploring Gender Differences in Fatwa through Machine Learning. *Journal of Cultural Analytics*, 9 (3).

DOI: <https://doi.org/10.22148/001c.116368>

Publisher: McGill University, Department of Languages, Literatures, and Cultures

Version: Published Version

Downloaded from: <https://e-space.mmu.ac.uk/635193/>

Usage rights:  [Creative Commons: Attribution 4.0](https://creativecommons.org/licenses/by/4.0/)

Additional Information: This is an open access article published in *Journal of Cultural Analytics* by McGill University, Department of Languages, Literatures, and Cultures.

Enquiries:

If you have questions about this document, contact openresearch@mmu.ac.uk. Please include the URL of the record in e-space. If you believe that your, or a third party's rights have been compromised through this document please see our Take Down policy (available from <https://www.mmu.ac.uk/library/using-the-library/policies-and-guidelines>)

ARTICLE

Exploring Gender Differences in Fatwa through Machine Learning

Emad Mohamed¹, Raheem Sarwar²¹ Nazarbayev University, ² Manchester Metropolitan University

Keywords: fatwa analysis, gender and religion, machine learning, topic modeling, classification, regression

<https://doi.org/10.22148/001c.116368>

Journal of Cultural AnalyticsVol. 9, Issue 3, 2024

This paper focuses on exploring the differences in inquiries made by men and women within a religious context. Additionally, we aim to ascertain whether it's feasible to forecast the popularity of answers and the factors contributing to their popularity. To achieve this, we compile a new dataset comprising 40,000 question-answer pairs categorized by gender and popularity. These are sourced from online question-and-answer platforms. Our methodology involves comprehensive experimental analysis, utilizing advanced Arabic text preprocessing alongside machine learning algorithms. We concentrate on two primary objectives: predicting the gender of the questioner and forecasting the popularity of answers. Furthermore, we delve into thematic variations based on gender and address pivotal research queries that offer new perspectives within this domain. These include investigating the differences between questions posed by women versus men, exploring the potential for automated classification of queries by gender, predicting the popularity of fatwas, and identifying the contributing factors to their popularity. Our experimental findings demonstrate a 98% accuracy in gender prediction, precise predictions of popularity with minimal margin for error, and the identification of topics and their associations that are more inclined towards either men or women. We intend to share both the dataset and the source code openly with the research community.

1. Background and Introduction

Research on Muslim women has recently grown rapidly (Faiz et al.; Khan and Mollah; Kloos and Ismah). This may be due to shifting attention after women in the industrialized world gained considerable rights (Maftuhin; READ and BARTKOWSKI; Nikjoo et al.; Baboolal; Murrar et al.; Abu-Ras and Itzhaki-Braun). Most prominent among the issues of Muslim women are how (and why) Muslim women wear the hijab (Abu-Lughod; Acker; Brenner) and Muslim women's political participation (Finlay and Hopkins; Akbarzadeh and Roose; Bhimji). A recent study of Islam in the English Wikipedia has found that the tenth most salient collocate of the adjectives Muslim/Islamic is the noun woman, ahead of such common collocates as conquest, jurisprudence, state, art, philosophy, terrorism, fundamentalism, and even prophet (Mohamed), but it remains true that "studies on Muslim women's online activities remain few and far between" (Piela). Most of these studies focus on Muslim women's online activism and include their

attempts to challenge male dominance. Alexandros Sakellariou (Sakellariou) used discourse analysis tools to examine female Greek converts to Islam in their digital presence, using their conversion stories as a means of understanding their social milieu as well as their digital and non-digital identities. Rahman, Fung, and Yeo used a small corpus of 1480 online comments to study Canadian attitudes towards the Hijab (the Muslim women's headscarf). The authors made use of computational tools, namely sentiment analysis, as well as content analysis for their investigation (Rahman et al.). One common trait of these studies is that they used small samples and focused on specific issues (Al-Ghadir et al.; Al-Ghadir and Azmi; Al-Sarem and Emara).

In this paper, we introduce a new dataset containing 172,000 question-answer pairs written in Arabic from a religious question-answering platform and show its usefulness to answer important questions such as what are the differences between questions posed by men and women, what makes an answer popular? The answers to such questions add new insights to the existing knowledge in fields of social sciences, cultural analytics, personal law, economics, technology, education, medicine, religious studies, information retrieval, or digital text forensics. Although we limit our investigation to the aforementioned questions, the novel dataset introduced in this paper can be used for a variety of purposes including (1) studying linguistic gender variation since the questions are marked for the gender of the questioner, (2) tracking a specific issue as the questions span 17 years, (3) tracking the rise and fall of specific themes, thanks to the view numbers associated with different dates, (4) Natural Language Processing tasks like summarisation and document similarity in a religious domain, and (5) tracking authority shift in the Muslim world.

Islamic Question Answering sites such as [IslamWeb](#) is similar to other community question-answering platforms by not only allowing people to post a wide variety of questions but also by offering qualified legal scholars a chance to browse and answer any of them. Typically, in these services, new questions can be formulated at any moment, and receive several responses from different qualified legal scholars (Figueroa)

To describe our dataset we first need to illustrate the concept of Fatwa in Islam. Britannica provides an accurate definition of fatwa (*Fatwa*):

" Fatwa, in Islam, is a formal ruling or interpretation on a point of Islamic law given by a qualified legal scholar (known as a mufti). Fatwas are usually issued in response to questions from individuals or Islamic courts. Though considered authoritative, fatwas are generally not treated as binding judgments; a requester who finds a fatwa unconvincing is permitted to seek another opinion. "

While usually directed to scholars of Islam, fatwas span a wide range of topics including politics, financial matters, family problems, and even medical issues, among many others (Agrama; Ismail and Baharuddin; Dahlan et al.; Adel and Numan). Traditionally, fatwas were issued by institutions and institutional scholars, but now most fatwas are issued online, either by individual scholars or on web portals belonging to religious institutions. Two main websites (fatwa portals) dominate the fatwa market:

- [IslamWeb](#) is a comprehensive website with fatwas, articles, videos, and many other features for Muslims. According to [SimilarWeb](#), it ranks 3,927 globally and is visited by 20.23 million viewers a day with the top countries being Egypt, Saudi Arabia, Algeria, Morocco, and France. The website is owned by the Qatari Ministry of Religious Affairs. It is worth investigating why this website is so popular among Muslims from so many Muslim countries through social science research methods. This study is beyond the scope of this article.
- [Islam Questions and Answers](#) has a global rank of 6,181 and has 13.66 million daily visits. The website was founded and run by the Syrian/Saudi scholar Muhammad Salih Al-Munajjid. The top countries visiting IslamQA are Saudi Arabia, Egypt, the United States, the United Kingdom, and the United Arab Emirates.¹

We attribute the popularity of these websites partially to the fact that they handle many questions as they answer around 200 questions every day and publish sensitive material while maintaining questioner anonymity. The anonymity guarantees that questioners can experiment freely, talk about their mistakes without being embarrassed, address wrongs without facing the consequences, and be judged based on the presented facts and not on external information (Jordan). A full investigation of the factors, social and otherwise, that belie the popularity of these websites is worthy of future research.

Fatwas raises many questions that experts in Islamic Studies and the sociology of Islam may find interesting. Our newly created dataset should thus be of interest to these scholars and others in computational humanities, digital humanities, computational linguistics, and computational social science.

Research Questions: In this paper, we present a new massive dataset and the transformations it went through. In addition, we showcase the usefulness of the dataset in answering three pertinent questions, which add new insights to the existing knowledge:

¹ The numbers were obtained on 15 April 2021, and they reflect the immediate period preceding this date. The numbers get updated regularly, so the readers may not get the exact numbers reported here.

1. What, if any, are the differences between the questions posed by women and those asked by men?
2. Can we automatically classify the questions as either male or female?
3. Can we predict fatwa popularity? And what makes a fatwa popular?

The questions follow a logical sequence in which question one focuses on examining existing gendered data, question two applies classification when no gender labels exist, thus discovering more questions asked by women and men, and question three assigns a measure of importance to the questions to which gender has been assigned. The questions thus seek to maximize the benefits of this dataset to the researchers in the fields of Islamic studies, women's studies, and gender studies.

The rest of this paper goes as follows: in section 2, we describe the dataset; in section 3 we describe preprocessing and the methods we use to answer the questions; in section 4 we present the results with a discussion of important issues; and in section 5, we conclude the study and outline some future research.

2. Dataset

The data used in this study is extracted from the fatwa portal [Islam Web](#), which is one of the largest fatwa portals online. According to the [SimilarWeb](#), IslamWeb has 17.97M daily visits, with a global rank of 3,897, and it ranks first in the category of *Community and Society > Faith and Beliefs*. Islam Web publishes 200 fatwas a day on average on a variety of topics with questions coming from around the world, but the portal does not publish any demographic data, which makes it hard to sort the questions and the answers based on the gender of the questioner. Since gender is one of the major demographic indicators and is of main interest to the authors, we have utilized the morphological nature of Arabic in search of gendered linguistic inflexion. Fortunately, Arabic, the main language of Islam Web, is a morphologically rich language inflected for gender, so in some cases, when the questioner speaks in a personal way, or when the scholar addresses the questioner personally, it becomes possible to identify the gender of the questioner. We use linguistic cues to extract the gender-identifiable fatwas from the fatwa collection.

2.1. Linguistic cues for gender identification

With Arabic being a grammatical gender language, speakers use gendered pronouns, nouns, verbs, and adjectives to refer to themselves and others. The word for doctor in Arabic is either *tabib* (male doctor) or *tabiba* (female doctor), so if a question has the expression *ana tabib*, this indicates the questioner is male. We also use the answers for the same purpose. *kama*

ta'alm means *as you (male) know*, while the female version is *kama ta'lamin*. We use these linguistic cues to assign gender to questions. We have a four-step approach as follows:

- starts with a seed list of gendered expressions, including pronouns, nouns, adjectives, and verbs. The list also includes regular expressions in the form *I gendered_noun* and *you gendered_verb*.
- extracts all the fatwas that are either male or female.
- check the intersection of the male and female fatwas and remove any common fatwas from the sets
- examine five hundred fatwas manually to check how accurate the method is.

After applying the four steps, we found 40458 with uncontested gender information out of 172,000 fatwas. The manual checks found no incorrectly assigned fatwas.

2.2. The resulting dataset

The resulting dataset comprises 40458 questions with 17221 asked by women and 23237 asked by men. Questions vary considerably in length with the average number of words per question at 116.83, with a standard deviation of 120.25, a median of 80, a minimum of 3, and a maximum of 3019. There are differences between the lengths of questions by men and women. For men, the average is 109.5 words (std = 117, min = 3, max = 3019, median = 74) while for women the average is 126.79 (std = 123.9, min = 4, max = 2109, median = 89). The dataset is distributed in a JSON streaming file with the keys: *question*, *answer*, *date* indicating the date of publication, *classification*, which is a coarse classification of the question provided by the website and *views*, which is the number of times the fatwa has been viewed. The dataset (40458 questions) has been randomly divided into a training set (85%, 34389 questions) and a test set (15%, 6069 questions). 10% of the training set has also been dedicated as a development set. This division is the same across all the experiments below. All the dataset metadata is available, directly or indirectly in the html. In [Figure 1](#), we have a screenshot of a recent fatwa with numbers marking the different pieces of the annotation:

2.2.1. QUESTIONS AND ANSWERS

The dataset includes questions by the public and their answers by Muslim scholars on various aspects of Islam. The average length of questions is 116.83 words, with a standard deviation of 120.25, a minimum of 3, and a maximum of 3019 (and a median of 80). The average length of answers is 223.3, with a std of 163.8, a min of 0, and a max of 5809 (median = 185). There does not seem to be a strong correlation between question length and answer length as the Pearson correlation is 0.21. In the current work, we use the questions and

الرئيسية « فقه العبادات » الرعية والأمانة « أحكام الرعية والأمانة »

حكم التصرف بأدوات العمل للاستعمال الشخصي

رقم الفتوى: 5763

328

0 0 0

Like 0

Tweet Share

تاريخ النشر: الإثنين 16 شعبان 1442 هـ - 29-3-2021 م

التقييم: ★★★★★

السؤال

بحكم عملي أقوم بطلب كل ما يحتاجه المكتب من الألف إلى الياء، وأنا المسؤولة عن التوزيع.
السؤال: في بعض المرات يطلب المسؤول، أو المشرف، أو شخص عادي متني أشياء مثل: المطهر، أو الكفوف، أو أقلام؛ لاستعماله الشخصي، يستعمله في البيت، وليس في العمل.
وأنا أيضا لما رأيت المسؤول يأخذ -وهو راتبه كبير جدا- فقتت بأخذ بعض الأشياء.
فما الحكم في ذلك؟ علما أن عندنا ميزانية للقسم، وكل سنة تجدد.

الإجابة

الحمد لله، والصلاة والسلام على رسول الله، وعلى آله، وصحبه، أما بعد:

فيجب التقيد في صرف هذه الميزانية بلوائح العمل، بغض النظر عن موقع الموظف، وغناه، وفقره! فهذه وكالة وأمانة يجب حفظها، وأداؤها على وجهها.
ولا يخفى أن شراء المستلزمات الشخصية التي لا علاقة لها بعمل المكتب، من ميزانية القسم: [إنها] هو تمدد وتضييع للأمانة، لا يجوز فعله، وتجب التوبة منه.

وانظري للفائدة الفتوى: [5763](#).

والله أعلم.

Figure 1. A screenshot of a fatwa and its information: [1] is a categorization of the fatwa, and in this specific case it is about the rules on honesty, [2] is a title that reads the rule on the personal use of work tools, [3] is the fatwa ID number, [4] is the date of publication (up) and how many stars the fatwa received (down), [5] is the number of views, [6] is the question, and [7] is the answer. The red underlines indicate morphological information marking the gender of the questioner (female) and the blue underline is a link to another relevant fatwa.

the answers to predict both gender and views based on textual data, but the questions and answers could be used to study various issues in Islamic studies and the sociology of Islam. Linguistic gender variation is also a potential research question that is made possible by this dataset.

2.2.2. TITLES

A title is usually a summary of the questions styled as a journalistic headline and is meant to attract attention. Although we do not use titles in this paper, mainly because they don't have the gender cues and because their content is included in the body of the Fatwa, they hold quite some potential as they can be used in Natural Language Processing research for summarisation and title generation. They can also be used in text similarity experiments and in short text classification.

2.2.3. VIEWS

The website records the HTML hits each page receives. This is extremely useful information for fatwa popularity as it could tell us how many people are interested in this specific question, and with some abstraction, how important the theme or category of the fatwa is. There is a strong, but not perfect, correlation between the number of questions in a category and the total number of views per category (Pearson $r = 0.79$).

The questions and answers in the gender-labeled data have been viewed 180,738,210 times, with an average of 4465 per question, but this is not evenly distributed as the standard deviation is 10124.7, with a minimum of 12 and a maximum of 441706. The median is 2485, which indicates a right-skewed distribution. The 5 top-viewed questions in the corpus are about prayer (441706 views), masturbation (432674), concubinage (318502), sexual excitation and genital cleanliness (314662), and divorce (304675).

The counts above are based on the views as they were recorded on 8 February 2016, but the data also includes the views as recorded on 8 July 2020. The 2020 counts are available for only 34714 fatwas in the corpus since the website deletes fatwas from time to time. It is not clear why the website deletes some Fatwas. According to the 2020 counts, the 34714 fatwas have been viewed 371,142,364 times, with an average of 10691 views per fatwa. We use views in a regression experiment in which we try to predict the number of views based on the textual content of the question. The purpose of prediction is two-fold: (i) it is used in the regression model to explain the contribution of each theme to the number of views, which helps rank those themes in terms of importance, and (ii) the model can be used to predict how popular a new question may become in future.

2.2.4. CATEGORIES

Each fatwa is categorized according to a hierarchical set of labels. For example, a question on whether it is Islamically legitimate to work on improving the Arabic Wikipedia is assigned the label **Main** – > **Thought, Politics and Art** – > **Culture and Thought**² while a question from a young woman complaining against her father who does not let her drive her own car is classified as **Main** – > **Family Matters** – > **Women’s Issues**.³ These labels are useful in obtaining a coarse-grained idea of the range of issues raised on these platforms. This could be used in Natural Language Processing and Machine Learning research for learning (hierarchical) classifications.

2.2.5. TEXTUAL EVIDENCE

When the muftis provide answers, they usually support their answers with textual evidence from the Qur’an, the Prophetic traditions, quotes by prominent scholars, or scientific research. These could be useful in understanding the sources governing Muslim thinking. They could also be used in machine learning and digital humanities research for intertextuality detection and text reuse.

² <https://www.islamweb.net/ar/fatwa/385525/>

³ <https://www.islamweb.net/ar/fatwa/430699/>

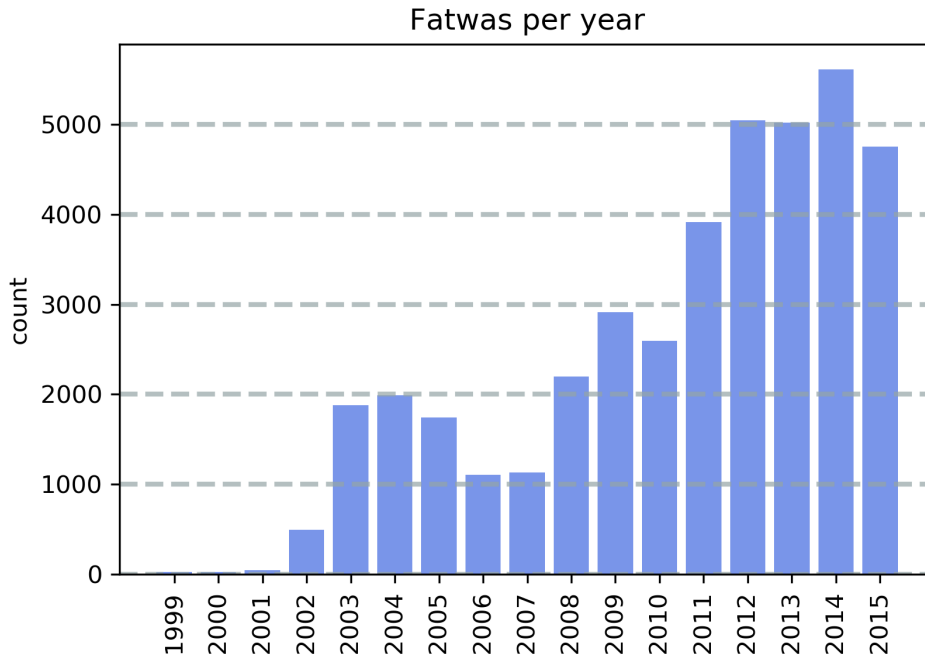


Figure 2. Fatwa counts per year from 1999 to 2015

The resulting dataset, as shown in [Figure 2](#) spans 17 years, from 1999 to 2015, with the number of views per fatwa available for February 2016 and July 2020, which could be used for both tracking fatwa popularity through time and regression analysis.

3. Methodology

To answer the first question: *What, if any, are the differences between the questions posed by women and those asked by men?*, we use topic modelling to find the themes of the questions raised by men and women and the odds ratios to determine which themes are more female and which tend to be more male. To answer the second question: *Can we automatically classify the questions as either male or female?*, we use text classification, mainly through automatic machine learning, and we also try to find which lexical items are more associated with men and women. For the third question: *Can we predict fatwa popularity? and what makes a fatwa popular?*, we use text regression using two popular algorithms: linear regression and random forests regression combined with topic modelling, which is used for explanation. As Arabic is a morphologically rich language, we use stemming throughout.

3.1. Stemming

A white space-delimited unit in Arabic is usually made up of zero or more prefixes, a lexical item, and zero or more suffixes. Common prefixes are conjunctions, prepositions, and the definite article while common suffixes are possessive and object pronouns. For example, the orthographic unit *fsykfykhm*, depicted in [Figure 3](#), is made up of the conjunction *f*, the future

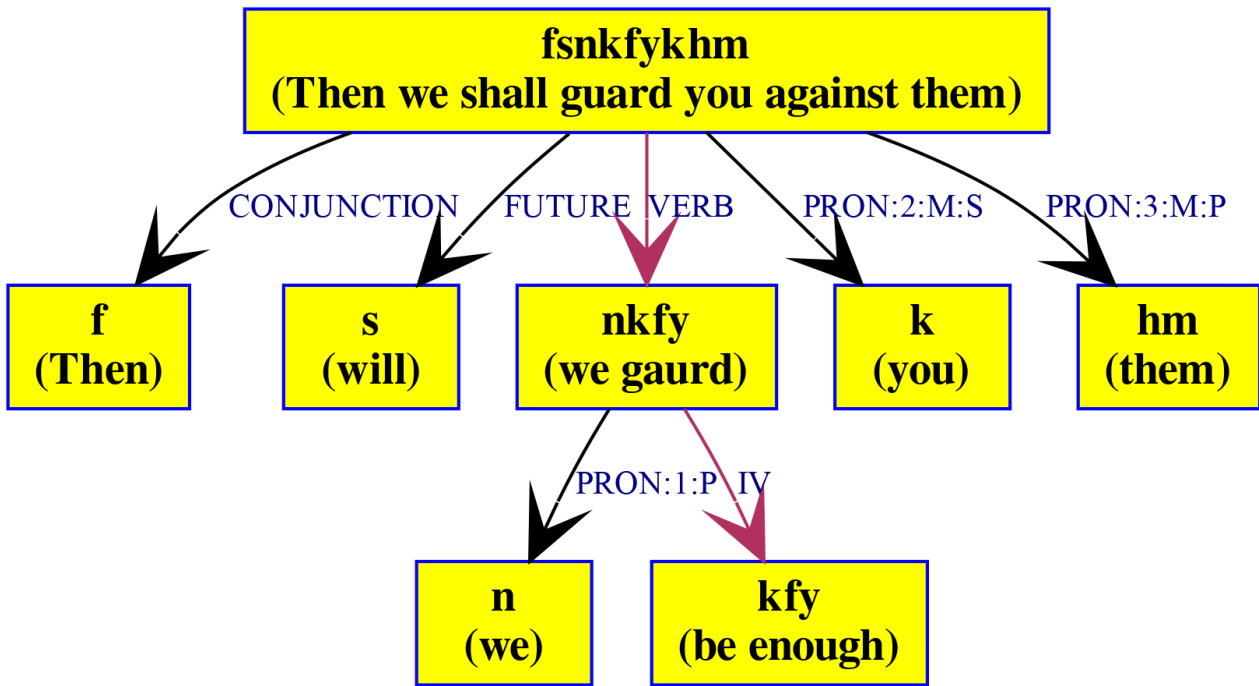


Figure 3. The morphological structure of the Arabic orthographic unit *fsnkfykhm*. The path to the stem is marked by maroon-coloured arrows.

particle *s*, the verb *ykyfy*, the direct object *k* and the indirect object *hm*. The verb itself is made up of the prefix *n* and the verb stem *kfy*. In most of what we do in this paper, we are mostly concerned with the stem.

For stemming to happen, the text is first run through morphological segmentation, which sets segment boundaries within words. The segments are then passed through a part of the speech tagger that assigns grammatical tags (e.g. NOUN, VERB, ADJECTIVE) to these segments. The stem is the main lexical unit in the word and must have a lexical tag (NOUN, VERB, ADJECTIVE). All other tags are discarded in the topic modeling experiments while in lexical experiments they are retained since they are useful for style differentiation. Stemming is achieved through the use of *ArabicSOS*, which is specialized in segmentation, orthographic standardization, and stemming of classical and religious Arabic (Mohamed and Sayyed). The stemming effect on the dataset is dramatic as the number of words in the questions is 4729279 with 168827 unique words. The number of segments is 7868382 with 48553 unique segments. The unique segments are hardly 29% of the original unique words.

3.2. Topic modeling

Topic modelling is a way of summarising text documents into collections of thematically related words called topics. Due to the nature of Arabic, there is a high type-token ratio as there are too many unique words, and running topic modelling with this is not very useful. For this reason, we use the stems as input to the topic modelling software. We use Mallet (McCallum) for

running the topic modelling and we are interested in two Mallet outputs: (1) the keys of the topics, or the lexical items constituting the topic, and (2) the probabilities associated with the topics, especially the probability of each topic in each document as we use these probabilities as input to the machine learning classifiers below. For consistency in topic modelling, we use the same 50 topics for all the experiments below.

3.2.I. CLASSIFICATION

For gender prediction, we use supervised machine learning (Sarwar and Mohamed). There are two main settings in these classification experiments:

- *text classification*, in which the dependent variable is the gender (male or female) and the predictive variables are made up of the textual content of the fatwa. The textual content could be word unigrams (each variable is a single word) or word bigrams (each variable is either a unigram or a bigram. For example, in the sentence: *The sky is blue*, the unigrams are ['the', 'sky', 'is', 'blue'] and the bigrams are ['the sky', 'sky is', 'is blue']). The values of these variables will be the frequencies of these n-grams in each document. We do not use raw frequency, but we use **TFIDF** (Term Frequency Inverse Document Frequency), which computes the most distinctive lexical items for each document and is thus more conducive to accurate classification.
- *topic classification*, in which we use the output of topic modelling for classification. When we run topic modelling, we obtain, for each document in the corpus, the probabilities of each topic being in that document. If we use 50 topics, then, for each document, we have 50 probabilities corresponding to the 50 topics. These probabilities can then be used to predict the gender of the document. Refer to the section below for results.

3.3. Regression

The purpose of regression is to predict, for each question, how popular the question may be. Just like with classification, we use both the textual content and the topics for regression. In both regression and classification, we use a variety of algorithms that will be detailed below. For both regression and classification, we use 80% of the data for training and 20% for testing. We mainly use the *scikit-learn* (Pedregosa et al.; Mohamed and Sarwar) library, which has a wide range of algorithms.

4. Discussion of Results and Implications

In this section, we discuss the results of our experimental studies and their implications. Recall that, in this paper, we present a new massive dataset and the transformations it went through. In addition, we showcase the usefulness of the dataset in answering three pertinent questions, which add new insights to the existing knowledge.

4.1. Answer to Question 1

What, if any, are the differences between the questions posed by women and those asked by men?

In order to find out what the differences may be between the questions asked by women and those asked by men, we use topic modelling. We then use Odds Ratios (see Tables 3 and 4) to find out which topics are more correlated with men and which are more characteristic of women. Table 1 lists the 50 topics produced by Mallet ordered by their odds ratios. Since female-centric questions were assigned the zero class, odds ratios less than one are more associated with women while odds ratios larger than one are more characteristic of male-centric questions.

We can see from the table that there are clear differences between the questions asked by women and those raised by men. The top concern for women in this dataset is family matters (children, marriage, divorce). We can also see religious ritual questions concerning menstruation and fasting as menstruating women do not have to observe the obligatory fasting in the lunar Islamic month of Ramadan. The current study focuses on a convenience sample of Muslim men and women who are mostly in the Middle East, and the results may thus be hard to generalize. The interest in the family may echo women's interests in other cultures as well. A study that examined a corpus of happiness moments found out that women and men have different sources of well-being (Mohamed and Mostafa), as men's top source of happiness was related to games and sports while women's was more related to family and shopping.

While differences between Men and Women may not be news in general, the details of these differences within the religious domain are of interest. For example, Men seem to be interested in totally different things, such as ensuring that the job or work they have is conformant to the rules of Islam, working and residing in other countries, especially Western countries, and obtaining their residences and citizenship, marriage and divorce consequences, the differences among scholars of Islam, life, death, and the hereafter, the relationships between Muslims and non-Muslims, banking and usury, especially concerning getting a bank loan and whether this is permissible in Islam, trade and transactions and rules on how to make

Table 1. The 50 topics produced by Mallet topic modelling ranked by Odds Ratios from most female (topic 4) to most male (topic 47)

OR	Topic	Theme
0.0	4	husband married family home divorce children problems refusal right
0.0031	36	period menstruation
0.0627	49	tradition problem talk people hope return
0.0951	0	love heart feel person problem haram
0.1069	3	soul feel life fear problem Quran pray depression
0.1424	29	mother sister brother father uncle
0.1707	27	women dress hijab hair religious beard
0.2459	22	call tell family refuse to agree to give
0.2718	42	marriage man woman engage proposal refuse family religious decent
0.2727	10	obsessive_compulsive_disorder devil thoughts
0.2765	41	Ramadan fasting expiation feeding_the_poor
0.3315	13	hit talk anger in_front_of people problem treatment insult
0.352	18	internet social friend love telephone
0.3925	20	ablution urine secretions semen prayer
0.4183	19	magic sleep Quran jinn evil_eye
0.6063	38	money salary haram amount expenses help
0.6888	1	prayer subsistence bless
0.6954	16	illness doctor-patient hospital psychiatrist fetus
0.6993	14	university school student test diploma graduation
0.9859	5	oath vow Mushaf lie
1.0329	31	wash water ablution foot hair
1.3784	39	question thank_you answer bless_you
1.383	17	life problem big family try leave
1.4215	12	time years months week hour
1.4435	24	please answer question
1.4598	33	clothes water dirty wash urine bathroom floor
1.7441	28	marriage certificate dowry agent conditions
1.8853	43	reading Quran prayer miss
1.9005	44	divorce anger return home dispute
2.1441	26	home husband visit family
2.2113	37	haram software game song music film draw tv
2.3708	45	unbelief religion unbeliever insult mockery heart repent
2.4811	7	time prayer athan masjid congregation dawn noon sunset
2.9943	8	food drink wine alcohol pork restaurant
3.1749	21	pilgrimage
3.4045	34	zakat alms
3.4308	2	inheritance estate
4.1313	9	flat land house rent property
4.4059	35	married children another_marriage divorce daughter
4.7719	46	paradise hell judgment chastisement
5.5224	30	country travel city egypt saudi work holiday
5.9501	40	repent masturbation sex lust sin forgive haram
7.2141	32	opinion scholars question disagreement law evidence
7.484	15	car accident court rights compensation claim
7.9649	11	company employee service office salary government
8.4999	6	Islam religion young_man live Christian france language foreign
12.4295	25	sale company buy price shop dollar proceeds trade commodity commission

OR	Topic	Theme
14.3603	23	verse Quran book exegesis story chapter
19.339	48	bank loan usury installment interest haram
49.4263	47	masjid people congregation innovation sermon heresy

profit and the permissibility of commissions, and questions about praying in mosques vs. praying at home. For example, a representative fatwa for topic 48, the second most male topic reads:

I'm a 32-year-old young man. I suffer from weak memory (I'm very stupid), and I have failed to hold any job due to this memory problem. Is this a license for me to keep my money in an Islamic bank, knowing that the three Islamic banks in Egypt (Faisal, Al-Baraka and Abu Dhabi) deal in treasury bills? Thank you!⁴

While a representative fatwa for topic 4, the most female topic reads:

My husband is a drug addict, and we've had three children together, the youngest being six years old. He wants more kids, but I do not feel like it. I do not use contraceptives, but his addiction stands in the way. Am I committing a sin by not seeking to get pregnant? Please reply quickly. I do not know who else to ask. It's been months, and I cannot find a solution.⁵

The use of topic models also gives us the chance to examine which topics go together. [Figure 5](#) shows the relationships between the topics. To create this network, we only considered the most probable topic pairs. To illustrate, let's examine topic 4, the most common topic in the questions asked by women. When we have topic 4 as the most probable theme in a certain document, topics 13 and 26 are the most probable second topics in that document, so we connect them together. For topic 4, this results in the graph shown in [Figure 4](#). What this indicates is that if a question involves matters of problems with their husband, family, children, and divorce, then it is also very likely to mention that the husband hits the questioner, insults her in front of other people, is mad at her and treats her badly (topic 13), and that this may come in the context of a discussion about visiting family (Topic 26). Topic 4 is also very likely to be invoked in the context of topic 44 (divorce anger return home dispute), topic 38 (money salary haram amount expenses help), topic 49 (tradition problem talk people hope to return), and topic 17 (life problem big family tries to leave).

⁴ <https://www.islamweb.net/ar/fatwa/294250>

⁵ <https://www.islamweb.net/ar/fatwa/138500>




Figure 4. The topics connecting to and from Topic 4

4.2. Answer to Question 2: Predicting Gender

Can we automatically classify the questions as either male or female?

In our specific case, we want to predict the gender of the author of the question even though we may use the answer as a means to this prediction. If the answer is addressed to a female, based on the morphology, then it is certain that the question was asked by a female. If the addressee is male then the question was asked by a male. In the cases where the morphology does not help identify the gender of the questioner, we do not use that specific Fatwa in training the machine learning classifier.

We run several experiments to predict gender. In all cases, we use the Term Frequency Inverse Document Frequency (TF-IDF) for feature extraction and an n-gram range of (1, 2), which means we use both individual words and two consecutive words as features. we use three algorithms in this task: Logistic Regression, Random Forests, and Support Vector Machines. As far as the input text is concerned, we vary the input to be either the question alone, the answer alone, or a combination of the answer and the question.

[Table 2](#) lists the results of these experiments. Using both the questions and answers as input with the Support Vector Machines algorithms yields the best results with precision, recall, and an F1 score of 0.98, which is a very high number, and means that we can predict the gender of the one who asked the question with 98% score. Using topic probabilities in gender classification did not yield good results, compared to lexical input, as the best result was an *F1* score of 0.68 using Logistic Regression.

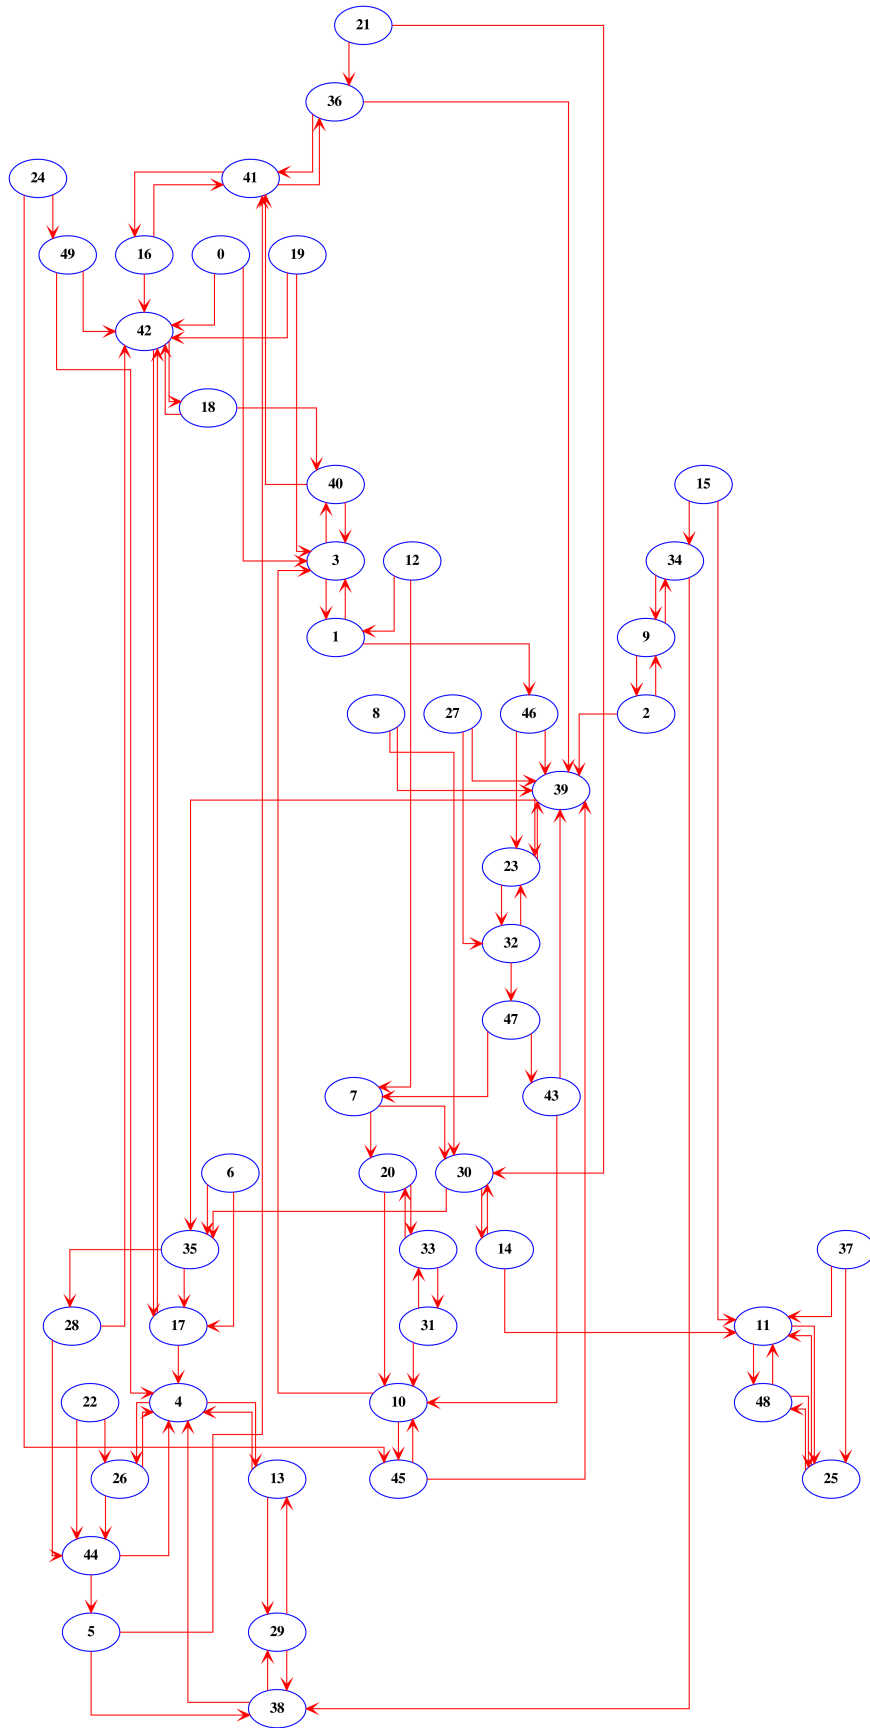


Figure 5. A network of the topics discovered by Mallet. Only the second most probable topic per document is connected to the main topic

Table 2. Predicting gender using the questions, answers, or a combination thereof.

Input	algorithm	Precision	Recall	F1
Questions	Logistic Regression	0.84	0.83	0.83
	Random Forests	0.81	0.80	0.79
	Support Vector Machines	0.82	0.82	0.82
	Logistic Regression, Topic models	0.71	0.68	0.69
Answers	Logistic Regression	0.94	0.94	0.94
	Random Forests	0.92	0.92	0.92
	Support Vector Machines	0.96	0.96	0.96
Questions + Answers	Logistic Regression	0.94	0.93	0.93
	Random Forests	0.92	0.91	0.91
	Support Vector Machines	0.98	0.98	0.98

While logistic regression is not the best-performing prediction algorithm, one major advantage of its use is its high interpretability. One can use the coefficients of the features (i.e. the lexical items) to see which words are more likely to be used by males and which are more often used by females. This also applies to the words and phrases more commonly used by the muftis when the questioner is male or female. This also helps with the possibility of generalization. We have extracted 40,000 out of 170,000 fatwas based on morphological clues. If these morphological clues distinguish male from female fatwas, then the possibility of an accurate classifier that generalizes beyond this small corpus is limited. If, on the other hand, the features are not limited to the morphological cues then we can use the classifier to predict much more than the small dataset. Examining the top 100 male lexical features and their equivalent female ones shows that the classifier is capable of generalization. The top 10 male features are not morphological.

4.3. Answer to question 3: fatwa popularity

Can we predict the popularity of a fatwa based on its content? To answer this question, we perform text regression where the predictor variables are the lexical items of the questions, and the response variable is the number of views each question receives.

To perform the regression, we convert the text into vectors. We do so through TFIDF as explained above. When we compute the number of views, we do so based on the number of views per day. Since the fatwas appeared online on different dates, it may not be fair to compare the raw views. This is the reason we divide the number of views by the number of days a fatwa is online, so we may obtain daily views. The daily views, rather than the raw ones, are what we want to predict. As shown in [Figure 2](#), the fatwa dates range between 1999 and 2015, but the views were recorded in February 2016 and July 2020. Since dates of publication are available for each fatwa, we subtract the publication date from 8/7/2020, the days the fatwa was downloaded, to obtain the number of days the fatwa was online. We then divide the number of views by the number of days to obtain the views per day. The

Table 3. Odds ratios for the most male lexical items. Male words are mostly generic while female words are predominately morphologically feminine.

Arabic	English	OR
إلا بإذن	except with permission	179
مُلزم	obligatory, committed	178
لا تعود	does not return	176
مقبل على	about to	168
مع بعض	with some, together	165
تتصح	advise	164
يجيء أحدهم	one of them comes	163
المالية	the financial	155
عملي هذا	this work of mine	148
هل أنا	Am I	147
وراجع الفتاوى	And review the fatwas	146
الباطل	falsehood	145
وأكثر من	And more than	144
لخطبتها	For her engagement	141
وخطيبتني	And my fiancee	141
كنت تقوم	you used to do	139
أنا متأكد	I am certain	134
من رأسه	of his own accord	131
للفتوى	for the fatwa	118
أن تبادر	to take the initiative	117

Table 4. Odds ratios for the most frequent female lexical items. Male words are mostly generic while female words are predominately morphologically feminine.

Arabic	English	OR
وانظري	and look	000
وراجعي	and review	000
انظري	look	000
أنا فتاة	I'm a young woman	000
واعلمي	and know that	000
السائلة	the questioner	000
وتكفيك	and it will suffice you	000
أنا أم	I'm a mother	000
تعديه	his aggression	000
تحسين	you feel	000
تأتيك	it comes to you	000
فاعلمي	then know	000
زوجي	my husband	000
تراضني	make (someone) happy	000
تعلمين	you know	000
أنا صاحبة	I am the one with	000
وزوجي	and my husband	000
ترضيه	you make him happy	000
تحسين الوضع	making things better	000
أنا امرأة	I'm a woman	000

Table 5. Regression results where the number of daily views is the dependent variable. Experiments report on using whole words vs segments, with unigrams or unigrams and bigrams, with the algorithm being either linear regression or Random Forests regression

Algorithm	Features	R^2	MAE
Linear regression (2016)	word unigrams + bigrams	0.99953	5.62
	word unigrams	0.99947	16.42
	segment unigrams	0.8675	52.34
	segment unigrams + bigrams	0.9995	7.83
Random Forests (2016)	word unigrams + bigrams	0.8647	3.49
	word unigrams	0.8516	3.733
	segment unigrams	0.8529	3.82
	segment unigrams + bigrams	0.863	3.61
Linear Regression (2020)	word unigrams + bigrams	0.999	5.584
	word unigrams	0.999	15.58
Random Forests (2020)	word unigrams	0.85	2.9
	unigrams + bigrams	0.86	2.67
	topic probabilities	0.87	3.1

views per day is the number we predict using the various regression models. We use different algorithms and compare the results to see which one best predicts the daily views per question. We also examine the feature ranking for features that affect the regression model. In terms of evaluation, we evaluate the regression model based on two things: (1) the R^2 value and (2) how good it is at predicting the views in a test set based on the Mean Absolute Error (MAE). The R^2 is a measure of goodness-of-fit and is usually interpreted as how much variation is explained by the independent variables. An R^2 of 0.6, for example, indicates that 60% of the variation in the dependent variable can be explained by the independent variables. The R^2 is known to be sensitive to the number of independent variables as it increases proportionally with the increase in the number of independent variables. This is a problem in text regression as the number of independent variables is very large, which may lead to inflated R^2 values. For this reason, we also use prediction to evaluate the regression model. The MAE is a measure of the average distance between the actual values and the values predicted by the model and can thus be useful in comparing various models.

The results of regression, as shown in [Table 5](#) indicate that when we use a combination of word unigrams and bigrams, both linear regression and Random Forests regression do a good job predicting views per day given the question as input. Linear regression has an R^2 of 0.99953, which is near perfect but may also indicate over-fitting for the training data. The mean absolute error is 5.62 on the test set. This is a good MAE given that the standard deviation of the test set values is 9.55. Random Forests do an even better job at prediction, probably due to their non-over-fitting, with an MAE of 3.49. This indicates that predicting the views per day based on the textual content is feasible. The experiments also show that segmentation does not help as every whole-word experiment yields better results than its segmented

Table 6. The topics contributing the most to fatwa popularity.

Rank	Topic	Keywords
1	42	marriage, man, woman, engage, proposal, refuse, family, religious, decent
2	40	repent, masturbation, sex, lust, sin, forgive, haram
3	20	ablution, urine, secretions, semen, prayer
4	41	Ramadan, fasting, expiation, feeding (the poor)
5	25	sale, company, buy, price, shop, dollar, proceeds, trade, commodity, commission

counterpart. We attribute this to the availability of data. With our large data set, there is no need for segmentation, whose main purpose is to combat data sparseness.

The Random Forests algorithm is superior to the linear regression one in prediction. One other nice facet of RF is that it also produces a list of the top features used in regression. Perhaps one should examine what features in the input are responsible for these views. For this, we use the most important features (lexical n-grams) as produced by the RF algorithm. An examination of the top 100 features in the unigram +bigram model reveals that the most important concepts that trigger views are related to sex, cleanliness, and prayer. For example, the top 10 grams can be translated into: *through the condom, is foreplay permissible, caressing the, anal, condom, inserting part of, intercourse, prayer, at the right time, masturbation, and orgasm.*

Perhaps a more effective way for finding out popularity triggers is to find what themes are more viewed. For this purpose, we use the probabilities of the topics produced by topic modelling as features and the views per day as the dependent variable, to rank the topics in terms of their importance to the Random Forest regression algorithm. This ranking shows how much each topic contributes to the number of views per day. [Table 6](#) lists the five most important topics, and we can see that questions about marriage and engagement top the list. A common theme here is the problem of the family refusing the prospective fiance(e), and the question is usually what one should do in such a case. This is followed by the question of (premarital and extra-marital) sexual activity and how one should repent. We then have topics on cleanliness and prayer, fasting the month of Ramadan, and what one should do if they miss days of required fasting, and then we have financial questions most of which are about whether some financial transaction is halal (permissible in Islam) or haram (impermissible according to Islamic law). The prominence of sexual topics may be related to the fact that Islam is very restrictive in terms of premarital sex (Adamczyk and Hayes).

5. Conclusion

In this paper, we have introduced a versatile dataset with textual content and metadata that make it suitable for research in computational linguistics, Islamic Studies, and Computational Social Science. We have also shown

use case examples of the dataset about questions of thematic analysis, text classification, and text regression. The dataset, with the use cases provided, has great potential for further formal and content-based research.

We found that there are clear differences between the questions asked by women and those raised by men. The top concern for women in this dataset is family matters (children, marriage, divorce). We can also see religious ritual questions concerning menstruation and fasting as menstruating women do not have to observe the obligatory fasting in the lunar Islamic month of Ramadan. We also examined the top 100 male lexical features and their equivalent female ones to show that the classifier is capable of generalization. The top 10 male features are not morphological. Our experimental findings demonstrate a 98% accuracy in gender prediction, precise predictions of popularity with minimal margin for error, and the identification of topics and their associations that are more inclined towards either men or women.

In the future, we are planning to investigate further questions including: (1) How do the muftis formulate their answers, and do they speak differently to men and women? (2) What is the authority frame that the muftis use to convince the reader of the soundness of their answers? and (3) What are the linguistic differences between male and female-centric questions? The dataset is ripe for investigation and we hope other researchers will find it useful for their research.

Data repository: <https://doi.org/10.7910/DVN/ASAJ4Y>

Submitted: January 08, 2024 EDT, Accepted: February 12, 2024 EDT



This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CCBY-4.0). View this license's legal deed at <http://creativecommons.org/licenses/by/4.0> and legal code at <http://creativecommons.org/licenses/by/4.0/legalcode> for more information.

WORKS CITED

- Abu-Lughod, Lila. "Veiled Sentiments: Honor." *Modesty, and Poetry in a Bedouin Society* (Berkeley: University of California Press, in Press) *Abu-Lughod Veiled Sentiments: Honor, Modesty, and Poetry in a Bedouin Society*, 1986.
- Abu-Ras, Ruba, and Yael Itzhaki-Braun. "The Complex Role of the Community in the Determination of Well-Being and Hope among Divorced Muslim Women." *Journal of Community Psychology*, 2023.
- Acker, Joan. "Hierarchies, Jobs, Bodies: A Theory of Gendered Organizations." *Gender & Society*, vol. 4, no. 2, June 1990, pp. 139–58, <https://doi.org/10.1177/089124390004002002>.
- Adamczyk, Amy, and Brittany E. Hayes. "Religion and Sexual Behaviors: Understanding the Influence of Islamic Cultures and Religious Affiliation for Explaining Sex Outside of Marriage." *American Sociological Review*, vol. 77, no. 5, 2012, pp. 723–46.
- Adel, Samiullah, and Muhammad Numan. "Online Fatwas in Pakistan Using Social Networking Platforms." *Ulumuna*, vol. 27, no. 1, June 2023, pp. 201–26, <https://doi.org/10.20414/ujs.v27i1.689>.
- Agrama, Hussein Ali. "Ethics, Tradition, Authority: Toward an Anthropology of the Fatwa." *American Ethnologist*, vol. 37, no. 1, Jan. 2010, pp. 2–18, <https://doi.org/10.1111/j.1548-1425.2010.01238.x>.
- Akbarzadeh, Shahram, and Joshua M. Roose. "Muslims, Multiculturalism and the Question of the Silent Majority." *Journal of Muslim Minority Affairs*, vol. 31, no. 3, Sept. 2011, pp. 309–25, <https://doi.org/10.1080/13602004.2011.599540>.
- Al-Ghadir, Abdul Rahman I., et al. "Gender Inference for Arabic Language in Social Media." *Discrimination and Diversity*, pp. 811–21, <https://doi.org/10.4018/978-1-5225-1933-1.ch037>.
- Al-Ghadir, Abdulrahman I., and Aqil M. Azmi. "A Study of Arabic Social Media Users—Posting Behavior and Author's Gender Prediction." *Cognitive Computation*, vol. 11, 2019, pp. 71–86.
- Al-Sarem, Mohammed, and Abdel-Hamid Emara. "Analysis the Arabic Authorship Attribution Using Machine Learning Methods: Application on Islamic Fatwā." *Advances in Intelligent Systems and Computing*, Sept. 2018, pp. 221–29, https://doi.org/10.1007/978-3-319-99007-1_21.
- Baboolal, Aneesa A. "(Under) Cover and Uncovered: Muslim Women's Resistance to Islamophobic Violence." *Victims & Offenders*, 2023, pp. 1–21.
- Bhimji, Fazila. *British Asian Muslim Women, Multiple Spatialities and Cosmopolitanism*. Palgrave Macmillan UK, 2012, <https://doi.org/10.1057/9781137013873>.
- Brenner, Suzanne. "Reconstructing Self and Society: Javanese Muslim Women and 'the Veil.'" *American Ethnologist*, vol. 23, no. 4, Nov. 1996, pp. 673–97, <https://doi.org/10.1525/ae.1996.23.4.02a00010>.
- Dahlan, Abdurrahman, et al. "Al-Buti's Thoughts on Maslāhah and Its Application in the Fatwa of World Fatwa Institutions." *Samarab: Jurnal Hukum Keluarga Dan Hukum Islam*, vol. 7, no. 2, 2023, pp. 1148–70.
- Faiz, Muhammad Fauzinudin, et al. "Challenging the Status Quo: Khaled m. Abou El Fadl's Perspectives on Islamic Legal Authority and the Restrictive Fatwa on Women's Solo Travel." *JIL: Journal of Islamic Law*, vol. 4, no. 1, Feb. 2023, pp. 47–66, <https://doi.org/10.24260/jil.v4i1.1071>.
- Fatwa*. <https://www.britannica.com/topic/fatwa>.

- Figueroa, Alejandro. "Male or Female: What Traits Characterize Questions Prompted by Each Gender in Community Question Answering?" *Expert Systems with Applications*, vol. 90, Dec. 2017, pp. 405–13, <https://doi.org/10.1016/j.eswa.2017.08.037>.
- Finlay, Robin, and Peter Hopkins. "Young Muslim Women's Political Participation in Scotland: Exploring the Intersections of Gender, Religion, Class and Place." *Political Geography*, vol. 74, Oct. 2019, p. 102046, <https://doi.org/10.1016/j.polgeo.2019.102046>.
- Ismail, Abdul Manan, and Ahmad Syukran Baharuddin. "Moderation in Fatwas and Ijtihad: An Analysis of Fatwas Issued by the MKI Malaysia Concerning the Covid-19 Pandemic." *AHKAM: Jurnal Ilmu Syariah*, 2022.
- Jordan, Tim. "Does Online Anonymity Undermine the Sense of Personal Responsibility?" *Media, Culture & Society*, vol. 41, no. 4, 2019, pp. 572–77.
- Khan, Borhan Uddin, and Md Al Ifran Hossain Mollah. "Protection through Constitutional Guarantees: The Case of Women, Children, and Backward Sections of the People." *The Constitutional Law of Bangladesh*, 2023, pp. 213–28, https://doi.org/10.1007/978-981-99-2579-7_12.
- Kloos, David, and Nor Ismah. "Siting Islamic Feminism: The Indonesian Congress of Women Islamic Scholars and the Challenge of Challenging Patriarchal Authority." *History and Anthropology*, 2023, pp. 1–26.
- Maftuhin, Arif. "Islamic Law, Disability, and Women in Indonesia: The Cases of Nahdlatul Ulama and Muhammadiyah." *Journal of Disability & Religion*, vol. 28, no. 1, Sept. 2023, pp. 13–27, <https://doi.org/10.1080/23312521.2023.2255860>.
- McCallum, Andrew Kachites. *MALLET: A Machine Learning for Language Toolkit*. no. 3, 2002.
- Mohamed, Emad. "Jewish, Christian and Islamic in the English Wikipedia." *Online-Heidelberg Journal of Religions on the Internet*, vol. 11, 2016.
- Mohamed, Emad, and Sayed A. Mostafa. "Computing Happiness from Textual Data." *Stats*, vol. 2, no. 3, July 2019, pp. 347–70, <https://doi.org/10.3390/stats2030025>.
- Mohamed, Emad, and Raheem Sarwar. "Linguistic Features Evaluation for Hadith Authenticity through Automatic Machine Learning." *Digital Scholarship in the Humanities*, vol. 37, no. 3, Nov. 2021, pp. 830–43, <https://doi.org/10.1093/llc/fqab092>.
- Mohamed, Emad, and Zeeshan Ali Sayyed. "Arabic-SOS: Segmentation, Stemming, and Orthography Standardization for Classical and Pre-Modern Standard Arabic." *Proceedings of the 3rd International Conference on Digital Access to Textual Cultural Heritage*, 2019, pp. 27–32.
- Murrar, Sohad, et al. "Predictors of Perceived Discrimination in Medical Settings among Muslim Women in the USA." *Journal of Racial and Ethnic Health Disparities*, vol. 11, no. 1, Jan. 2023, pp. 150–56, <https://doi.org/10.1007/s40615-022-01506-0>.
- Nikjoo, Adel, et al. "The Contribution of All-Women Tours to Well-Being in Middle-Aged Muslim Women." *Gender and Tourism Sustainability*, Jan. 2023, pp. 269–84, <https://doi.org/10.4324/9781003329541-17>.
- Pedregosa, F., et al. "Scikit-Learn: Machine Learning in Python." *Journal of Machine Learning Research*, vol. 12, 2011, pp. 2825–30.
- Piela, Anna. "Muslim Women Speak Online: Religion, Conversion, Activism, and Art." *Hawwa*, vol. 13, no. 3, Oct. 2015, pp. 271–78, <https://doi.org/10.1163/15692086-12341287>.
- Rahman, Osmud, et al. "Exploring the Meanings of Hijab through Online Comments in Canada." *Journal of Intercultural Communication Research*, vol. 45, no. 3, Apr. 2016, pp. 214–32, <https://doi.org/10.1080/17475759.2016.1171795>.

- READ, JEN'NAN GHAZAL, and JOHN P. BARTKOWSKI. "TO VEIL OR NOT TO VEIL?: A Case Study of Identity Negotiation among Muslim Women in Austin, Texas." *Gender & Society*, vol. 14, no. 3, 2000, pp. 395–417.
- Sakellariou, Alexandros. "Female Converts from Greek Orthodoxy to Islam and Their Digital Religious Identity." *Hawwa*, vol. 13, no. 3, Oct. 2015, pp. 422–39, <https://doi.org/10.1163/15692086-12341291>.
- Sarwar, Raheem, and Emad Mohamed. "Author Verification of *Nahj Al-Balagha*." *Digital Scholarship in the Humanities*, vol. 37, no. 4, Jan. 2022, pp. 1210–22, <https://doi.org/10.1093/lc/fqab103>.