Aalborg Universitet



Machine Learning and Asylum Adjudications: From Analysis of Variations to Outcome Predictions

Katsikouli, Panagiota; Byrne, William H.; Gammeltoft-Hansen, Thomas; Høgenhaug, Anna Højberg; Møller, Naja Holten; Nielsen, Trine Rask; Olsen, Henrik Palmer; Slaats, Tijs

Published in: **IEEE** Access

DOI (link to publication from Publisher): 10.1109/ACCESS.2022.3229053

Creative Commons License CC BY 4.0

Publication date: 2022

Document Version Publisher's PDF, also known as Version of record

Link to publication from Aalborg University

Citation for published version (APA):

Katsikouli, P., Byrne, W. H., Gammeltoft-Hansen, T., Høgenhaug, A. H., Møller, N. H., Nielsen, T. R., Olsen, H. P., & Slaats, T. (2022). Machine Learning and Asylum Adjudications: From Analysis of Variations to Outcome Predictions. *IEEE Access*, *10*, 130955-130967. Article 9984209. https://doi.org/10.1109/ACCESS.2022.3229053

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
 You may not further distribute the material or use it for any profit-making activity or commercial gain
 You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from vbn.aau.dk on: July 04, 2025



Received 17 November 2022, accepted 7 December 2022, date of publication 14 December 2022, date of current version 21 December 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3229053

APPLIED RESEARCH

Machine Learning and Asylum Adjudications: From Analysis of Variations to Outcome Predictions

PANAGIOTA KATSIKOULI[®]¹, WILLIAM H. BYRNE², THOMAS GAMMELTOFT-HANSEN², ANNA HØJBERG HØGENHAUG², NAJA HOLTEN MØLLER¹, TRINE RASK NIELSEN[®]¹, HENRIK PALMER OLSEN², AND TIJS SLAATS¹

¹Department of Computer Science (DIKU), UCPH, 2100 Copenhagen, Denmark ²Faculty of Law, UCPH, 2300 Copenhagen, Denmark

Corresponding author: Panagiota Katsikouli (pk@di.ku.dk)

This work was supported in part by the Project Data Science for Asylum Legal Landscaping (DATA4ALL).

ABSTRACT Individuals who demonstrate well-founded fears of persecution or face real risk of being subjected to torture, are eligible for asylum under Danish law. Decision outcomes, however, are often influenced by the subjective perceptions of the asylum applicant's credibility. Literature reports on correlations between asylum outcomes and various extra-legal factors. Artificial Intelligence has often been used to uncover such correlations and highlight the predictability of the asylum outcomes. In this work, we employ a dataset of asylum decisions in Denmark to study the variations in recognition rates, on the basis of several application features, such as the applicant's nationality, identified gender, religion etc. We use Machine Learning classifiers to assess the predictability of the cases' outcomes on the basis of such features. We find that depending on the classifier, and the considered features, different predictability outcomes arise. We highlight, therefore, the need to take such discrepancies into account, before drawing conclusions with regards to the causes of the outcomes' predictability.

INDEX TERMS Asylum adjudications, machine learning, automated decision-making.

I. INTRODUCTION

In Denmark, an individual who demonstrates a well-founded fear of being persecuted or faces a real risk of being subjected to torture is eligible for asylum. In Danish law, the exact legal thresholds reflect those established by international conventions, notably the 1951 Refugee Convention and the 1950 European Convention for Human Rights. These international treaties, however, remain largely silent when it comes to how states should go about assessing asylum claims. Asylum procedures are further subject to limited evidentiary material. As a result, national authorities are typically left to determine an individual's legal eligibility on a narrow basis consisting of an oral testimony, which may itself be hampered by several factors, including imprecise language

The associate editor coordinating the review of this manuscript and approving it for publication was Vicente Alarcon-Aquino¹⁰.

interpretation, insecurity or lacking trust towards the authorities among applicants, and psychosocial factors, such as PTSD, impairing the ability to precisely recount traumatic experiences.

The leaky ground, on which authorities must assess their subjective perceptions of asylum applicants' credibility, questions whether, in all cases, adjudicators make the correct decision. Moreover, the subjective element in these assessments raises questions on whether individual asylum cases could be afflicted by implicit biases or stereotyping amongst adjudicators. In fact, recent studies have uncovered significant correlations between decision outcomes and the experience and gender of the assigned judge, as well as correlations between asylum outcomes and entirely external events such as weather and political elections [1], [2], [3], [4], [5], [6], [7]. Various researches and technological tools have been used to mitigate for such discrepancies and contribute to analysing

legal sources and determining decisions outcomes, as a consultant in avoidance of biases. One emerging protagonist in this regard is Artificial Intelligence (AI) and in particular Machine Learning (ML).

ML brings two key advantages to the process: (1) it is highly automated and can efficiently exploit years of experience in historical data, and (2) mined rules can be applied deterministically, thereby outputting a predictable outcome, based on each case's characteristics. However, it also comes with a big caveat, widely known as the Garbage-In-Garbage-Out problem, where biases in historical decisions will be encoded in the data and therefore also become part of machine learned decision models [8], [9], [10].¹ To overcome this, the data on which a model is trained need to be reliable and representative, and we need to consider the risk of metrics and objective functions being shaped or even skewed towards political ends [11]. In addition as the data changes over the years and new variables become relevant, old models may become useless and thus need to be retrained on new, more relevant data [12].

In this paper, we analyse summaries of asylum-case decisions in Denmark, in order to explore variations in the recognition rates between different groups of applicants. Recognition rates refer to the fraction of received applications which were granted asylum. Our methodology can be summarized as follows. Firstly, we extract several features from the asylum applications in our dataset, we explore the varied recognition rates in relation to those features and discuss how such variations need careful interpretation on the basis of relevant laws. As we detail in Section II, literature plays host to a number of recent works that apply ML algorithms on datasets similar to ours, in an effort to use AI tools for uncovering biases and for building decision prediction models. Such studies often conclude on the predictability of the asylum decisions, and the existence of bias in the decision-making process. As a second task, therefore, we follow the paradigm and study the predictability of the decisions in our dataset. For this, we employ the various extracted features from the asylum applications in all possible combinations, and feed them as predictors to a number of classifiers. Different from similar works, our goal is to highlight the variation in accuracy results achieved by different classifiers and considered features. The conclusion of bias in a decision-making system requires deep understanding of the underlying system and its components, and it is influenced by the employed data and the used methodologies. As we show in Section III and discuss in Section V, different conclusions could be drawn from the results taken from different algorithms and data features. We wish to therefore impress on the importance of discussing the outcomes of different approaches, before drawing conclusions.

In summary, the paper's contributions are as follows:

- we present and analyse a large publicly available dataset of actual decision texts on asylum cases in Denmark,
- we present variations in the asylum decisions,
- we apply Machine Learning to study the predictability of the decision outcomes,
- we find that asylum decisions' predictability significantly varies with regards to the employed algorithms and data features, making it particularly important how one presents and interprets such results, before drawing conclusions,

II. RELATED WORKS

There is a growing body of literature both within law and data science addressing bias (or stereotyping) in legal decisionmaking. Recent scholarship has focused on offering empirical basis to investigate what drives a decision at law through computational methods. In this section we focus on literature that addresses: (a) bias in the outcomes of asylum cases, (b) the role of AI in decision-making and (c) computational approaches to law. For a broader literature review, see [13], [14], and [15].

A. BIAS IN ASYLUM DECISIONS

In [16], Chen and Eagle report biases stemming from several factors seemingly unrelated to the legal merits of the asylum applications themselves. The authors applied a random forest classifier to a dataset of 400,000 decisions from asylum hearings decided over a 32-year period and achieve an accuracy of 82% in outcome prediction, and suggest a number of extraneous factors which may indicate biases that can arise in asylum decision making. Dunn and Chen [17] use ML to analyse 600,000 asylum decisions of US judicial institutions, and achieve a predictability of 80% using only the identity of the judge and the applicant's nationality. Chen et al [2] have further found the prevalence of the gambler's fallacy in asylum adjudication, further indicating that decisions may be influenced by factors seemingly unrelated to the merits of a case. In [18] the authors analysed 400,000 asylum decisions and described the chance of receiving asylum in the US as a "refugee roulette". In particular, the rate of success is strongly affected by random assignment of a case to a particular immigration judge, but also gender and work experience of the judge prior to appointment, as well as the quality of the applicant's legal representation. In France, the SupraLegem project² applied ML on migrant expulsion decisions to achieve a prediction accuracy of 90-99% and finding statistics indicating bias amongst certain judges, as some appear to be more likely over time to reject appeals than others, in a case load where cases are distributed randomly between judges. Other types of biases, or decision factors, have been explored in [19], where the authors show bias against applicants in Canada persecuted

¹https://www.reuters.com/article/us-amazon-com-jobs-automationinsightidUSKCN1MK08G

https://hbr.org/2019/05/all-the-ways-hiring-algorithms-can-introduce-bias https://venturebeat.com/2020/09/24/in-facial-recognition-challenge-top-rankingalgorithms-show-bias-against-black-women/

²https://medium.com/@supralegem/the-impartiality-of-some-judgesundermined-by-artificial-intelligence-c54cac85c4c4

on account of bisexuality. In a study that examined decisions made by asylum officers and by immigration judges [20], from multiple countries, it was found that waning importance of human rights is more pronounced for asylum officers than for immigration judges after the attack on the World Trade Center in NY, USA. It is also reported that language heritage, specifically for asylum seekers from English, Spanish, and Arabic-speaking countries, substantially affects acceptance rates. The importance of implicit bias in immigration adjudication is highlighted in [21]. The article argues that the specific conditions under which immigration judges decide cases render them especially prone to the influence of implicit bias. Specifically, the article examines how factors such as immigration judges' lack of independence, limited opportunity for deliberate thinking, low motivation, and the low risk of judicial review allow implicit bias to drive decisionmaking. Other studies suggest that biases may arise in cases where asylum seekers are represented by immigration consultants as opposed to certified lawyers [22].

This body of literature gives insight into the existence of bias in asylum decisions. However, whilst many of these works follow a rigid methodology, there are instances where the methodology is insufficiently explained, or the results are not carefully validated. In particular, the accuracy metric of a ML algorithm is not sufficient to explain variations and conclude on the existence of biases. Careful investigation of the algorithm's performance is required, but also of the employed dataset and external information that is not accessible from the dataset (e.g., local or national laws.)

B. AI AND PUBLIC DECISION MAKING

There is a growing literature on automated decision-making in the legal domain and public administration [5], [13], [23], [24], [25], [26], [27]. Artificial intelligence itself has received controversial feedback from the community; on the one hand, machine learning algorithms have been criticised as unfair because of the bias and discrimination they may impose (see footnote examples in Section I). On the other hand, as we report further in this paragraph, machine learning has been presented as a valid solution for ameliorating the biases that may arise in human decision making. In [28] the suitability of AI is analysed for decision making in the legal profession, and is concluded that legal decisions are complicated and require a qualitative element which is not easily replaced by automation, but such tools may provide assistance in reaching those decisions. In [8] the authors propose that the use of AI can intensify asymmetries between the public authority and the immigrant. In [3] the authors engage in legal and ethical analysis of the use of AI in public administration and highlight the importance of ethical frameworks. In [29] the need for participatory design of AI decision-making systems in public administration is stressed. This view is echoed by [30] who stress the need to engage in multidisciplinary teams of researchers, policymakers and practitioners. Authors in [31] highlight how the use of AI for decision making in Canadian immigration applications may give rise to significant legal implications under international human rights law. Other studies propose the use of AI to overcome bias in decision making. In [32], the authors propose ML algorithms to predict extra-legal biases in judicial decisions. The need for in-depth empirical research to overcome uncertainties in algorithmic systems that arise in this context, is also proposed in [33]. The use of machine learning in empirical studies in the landscape of credibility is addressed in [34]. Machine learning for early acceptance or rejection of legal cases is proposed in [35]. In [36] the authors use natural language processing and ML for discovering violations of human rights law in relevant cases. The proposed models can predict the court's decisions with an accuracy of 79%. Machine learning for predicting outcomes is also applied in [37] where an accuracy of 75% is achieved in predicting decisions of the European Court of Human Rights. The authors show that predictive accuracy depends on continuous model updating to keep track with developments in the law (with accuracy dropping to 58% when there is delay in update.)

C. COMPUTATIONAL APPROACHES TO LAW

Computational law is a rapidly growing field of research. Movzina et al. [38] use argument based machine learning to learn the rules that have been used as arguments in legal decisions. The authors construct a tool through which machine learning can adapt to new rules as they arise in new cases and thereby further elucidate the development of legal rules. The authors of [6] use random forest classifiers to predict judicial voting behaviour. In [39], the authors use natural language processing, text similarity and machine learning to identify transpositions in the national implementation of EU legal derivatives. Text analysis has also been applied in order to analyse asylum/visa decisions for the purpose of determining the accuracy of systems for assessing the credibility of applications [40]. The use of blockchain systems in refugee systems is discussed in [41]. Complex networks and language processing are considered in [42] where citation network analysis and corpus linguistics are applied to the case law of European Courts [43]. Complex networks are also applied in [44] to the study of evolution of legal precedent.

Our work joins the group of research that applies natural language processing to extract features from texts of the decisions of asylum cases and ML to investigate the predictability of the decision outcomes, on the basis of the given application features. Unlike previous similar works, we conduct our research using decision extracts from asylum appeals cases published by the Danish Refugee Appeals Board. The accuracy of the employed algorithms on predicting the outcome of asylum decisions is critically assessed, in order to validate whether it could be attributed to existing biases or to artifacts of the dataset.

III. RECOGNITION RATES IN DECISIONS OF ASYLUM CASES IN DENMARK

In this section we present our employed dataset, which contains decision summaries from asylum cases at the appeal level in Denmark. Our dataset was taken from the publicly available repository Flygtningenævnets (FLN) Nævnsdatabase³ on the 20th of October 2020. It contains all summaries of decisions published until that day; that is, approximately 8,000 decisions to asylum applications tried by the Refugee Appeals Board, in the period 2003-2020. Those summaries were organized by the Secretariat of the Refugee Appeals Board's along three dimensions: year of decision, country of origin of applicant and type of asylum claim (e.g. "political conditions", "LGBT", "first country of asylum").

In order to get the data from the public repository, we extracted the html page containing all summaries across all three dimensions and then, by using the Python package Beautiful Soup, we removed all html markups and maintained the summaries' useful text.

This dataset comes with a number of peculiarities which drive our methodology and shape the way we interpret our findings. Firstly, the dataset is not necessarily representative of the full set of cases and decisions treated by the Refugee Appeals Board in Denmark, as not all decisions are made publicly available. However, the yearly recognition rate published by RAB coincides with the yearly recognition rate we calculated on our dataset, therefore, we can confirm that the employed dataset is statistically representative, at least with regards to the recognition rate. Although this first peculiarity is a generally known limitation [19] and the selection criteria for the published decisions are unknown to us, analysing them remains valuable as it highlights the variations in the decisions of such datasets, as well as their characteristics. Secondly, the dataset contains only asylum cases that have been rejected by the first instance decision making authority, the Danish Immigration Service. It is therefore not representative of the complete set of asylum cases received in Denmark. Thirdly, the dataset only contains the summaries of the decisions and not complete transcripts of the interviews with the applicants, nor information on the judges making the decisions.

Despite the full case files not being available, the considered summaries represent a relative rich source from which we can extract meaningful information about the asylum seekers' applications, the legal process followed and whether weight was accorded to external evidence. These summaries are furthermore particularly important since they represent the searchable database that legal practitioners have access to and rely on in terms of identifying relevant previous case law when preparing for a new decision. In our present analysis, we therefore focus on the considered features found in the summaries, and we concentrate on the impact such features might have on a case's decision being overturn. This choice further ties in with previous research conducted in other countries in terms of creating a comparable basis for which to compare across different legal jurisdictions (as we detail in the Related Works section). The question therefore asked in our work is: "Are there asylum applications with such characteristics that make the Refugee Appeals Board more likely to overturn the initial decision taken at the first instance?".

In the next paragraphs we describe the dataset, the application features we identified in the data, their distributions, as well as the techniques applied for their extraction from the decision text.

A. DATASET FEATURES

The extracted summaries are written in Danish. As such, most of the following procedures were adjusted for the Danish language. In order to identify features of interest and to estimate the extraction error, two Danish members of our group, one a legal expert, independently studied and manually extracted features from 50 randomly sampled cases. We concluded on the following feature-set: the applicant's country of origin (or nationality), the applicant's identified gender, the applicant's identified religion, the applicant's identified ethnicity, the year the applicant entered Denmark, their marital status, their involvement in political parties and organizations, military involvement or experience, whether the applicant has applied for asylum in another country before coming to Denmark, whether discrepancies were identified in the applicant's case, in cases of torture we check whether relevant investigation was carried out, the type(s) of asylum claim and finally, the Refugee Appeal's Board decision on the case.

Due to variations in spelling and to a non-standard format followed in the written text across all decisions, an extraction error is expected. For its estimation, we compared the data extracted manually to the data extracted automatically for these 50 cases and estimated the accuracy error. We found that the features Country of Origin, Year of Entry, Year of Decision, existence of divergences and the type(s) of asylum claim were extracted precisely by our automated system. For Gender, we found that in some instances, some cases concern whole families, but our automated system assigned those to a Male applicant. The accuracy for Gender was estimated at 0.96. Accuracy of at least 0.83 was achieved for the rest of the extracted features. The values are presented in Table 1. The distributions presented in Section III-B have been adjusted accordingly; in particular, we have adjusted the counts within each category by adding a value calculated by using the estimated error for the particular category. Next, we present the extracted features and their distributions, separately.

1) NATIONALITY

The country of origin of the applicant is explicitly stated in the FLN repository. The distribution of nationalities is shown in Fig. 1(a), ignoring (for readability) nationalities with less than 10 cases. Most applicants come from the Middle East and Somalia, explained by the wars in these regions in the last decade. A number of applicants do not wish to reveal their country of origin; that appears as "Unknown homeland".

³https://fln.dk/da/Praksis



FIGURE 1. (a) Fraction of asylum seekers per country of origin (nationality) (b) Acceptance rates per country of origin of asylum seeker (c) Fraction of asylum seekers per identified gender (d) Acceptance rates per identified gender of asylum seeker.

TABLE 1. Feature extraction accuracy.

Feature(s)	Accuracy
Entry Date, Date of Decision, Country of Origin	1
Type(s) of Asylum Claim, Divergences	1
Marital Status	0.97
Gender	0.96
Asylum elsewhere, Ethnicity	0.92
Torture Investigation	0.9
Political Involvement	0.84
Religion, Military Involvement	0.83

2) GENDER

We employed natural language processing, using the nltk library in python, to extract and count the presence of gender specific words, such as he/male/his/man vs she/female/her/woman. We note that for cases categorised as "LGBT" or/and "Gender Persecution" (360 cases in total), the extracted gender information is the one identified and stated in the application text. Figure 1(c) shows the distribution of identified genders. Evidently, the vast majority of cases concern male applicants. Although the age range of each applicant is not stated in the decision summaries, we note that, as per the analysed text, a small fraction of these cases concern minor applicants.

3) RELIGION

Even though an applicant's religion is not always explicitly stated in the texts, we could define regular expressions to extract the most common phrasing used for stating the applicant's religion. As an auxiliary tool we used the library spacy, where items from the tokenized cases were labeled in a number of categories, including religion, organization, nationality, etc. Whenever a religion instance was extracted for a case, we compared it against a long constructed list of available religions and belief systems, and accepted it if it appeared more than once in the dataset; else it would be dumped in a generic category "Others". For items not found in the religions list, we would count the frequency of the extracted item in the complete dataset. If that was more than 1 and less than 10, we dumped it in the generic category "Others". If it was more than 10, we assigned it its own category. Otherwise, the item was ignored. The religions distributions are shown in Figure 2(a). Atheists, agnostics and non-believers are merged in the category "Non-religious".

4) YEAR OF ENTRY

The year the applicant entered Denmark is usually not the same as the year when a decision is made by the Refugee Appeals Board. The average processing time for asylum applications at the first instance (Danish Immigration Service) is approximately 10-12 months. The case processing time at appeal level (Refugee Appeals Board) is slightly higher, approximately 14-18 months - and some cases may take significantly longer than those periods. What is more, the year the applicant entered Denmark is not necessarily the year the applicant made the asylum application, as in some cases, an applicant may have had a residence permit (e.g. familyrelated, work-related, or education-related) for a period of time before they decided to seek asylum. For its extraction from the text, we used regular expressions, as it is usually explicitly stated in the decision. Fig. 2(c) shows the number of analysed applications per year of entry in Denmark of the respective applicant. An obvious rise of incoming immigration appears within the last decade, especially at and around the year 2015, presumably related to the wars in Middle East and Somalia.

5) ETHNICITY

This feature was also extracted with the use of regular expressions. After we briefly inspected a large number of summaries manually, we identified common phrases used for presenting an applicant's ethnicity. Fig 3(a) shows the fractions of found ethnicities, where cases with less than 6 instances are omitted (for readability). We observe that in 1 out of 5 cases, the applicant's ethnicity either is not mentioned, or not detected by our method (possibly due to our regular expressions not being exhaustive).

6) BOARD's DECISION

Because our dataset concerns appeals cases, decisions are usually given as either upholding the Immigration Service's decision or not. To capture all possible ways of declaring the



FIGURE 2. (a) Fraction of asylum seekers per identified religion (b) Acceptance rates per identified religion of asylum seeker (c) Number of asylum seekers per year of entry in Denmark (d) Acceptance rates per year of entry of asylum seeker in Denmark.



FIGURE 3. (a) Fraction of asylum seekers per identified ethnicity (b) Acceptance rates per applicant's identified ethnicity (c) Fraction of decision outcomes (d) Fraction of asylum categories in the employed dataset.

Board's decision, we used nltk and regular expressions. Our approach treats cases where the main applicant is accompanying some minor and only the minor was offered asylum as granted asylum. The decision rates are presented in Fig. 3(c). The high number of rejected cases is consistent with the yearly recognition rates reported by RAB.

7) TYPES OF ASYLUM CLAIM

Next, we present the types of asylum claim assigned to the scraped cases by FLN. Each case is assigned at least one type. In Fig. 3(d) we present the disjoint types and the fraction of cases with corresponding labels. The dominating type is the generic "Agents of Persecution", followed by other generic categories, such as "General Conditions", "Private Law Relationships", "Political Relationships" and "Private Law Matters".

8) INVOLVEMENT IN POLITICAL PARTIES

We employed the nltk and spacy tools and organized the decision texts into 6-grams. The ones containing any found related entities (such as political parties and organizations) were isolated and further explored for inclusion/membership phrasing around the found entity. For example, if "Hamas" was identified, we extracted the 6-grams that contained the word and looked for phrases such as "member of", "part of" etc. Although we extracted exact information of the organizations/political parties an applicant might have been involved with, for the time being and the present analysis we have opted for a binary classification, that is, whether

an applicant has any (or not) involvement with organizations and political parties. Figure 4(a) shows that for the vast majority of applicants, such involvement was not detected in the decision text.

9) MILITARY INVOLVEMENT

Following a similar procedure, a binary classification of involvement in military groups is followed. The distribution found in our dataset is shown in Figure 4(c), where we can see that for most applicants such involvement was not detected in the decision text.

10) MARITAL STATUS

Using regular expressions, we mined the marital information of the applicants, which we grouped in four classes: married, not married, used to be married, and no marital information detected in the text. The distribution of these classes is shown in Figure 5(a). A large percentage of applicants is not married or does not disclose the relevant information. 40% of our dataset corresponds to married applicants.

11) DETECTED DIVERGENCES

By using regular expressions, we extracted information indicating whether divergences were noticed in the applicant's documents. For this, we constructed a list of possible ways that such divergences would be discussed in a legal context. As shown in Figure 5(c), a surprisingly large percentage of cases (nearly 40%) seems to have divergences detected in the asylum seeker's application.



FIGURE 4. (a) Fraction of asylum seekers for which involvement in political parties and organizations was detected (or not) in their case text (b) Acceptance rates depending on applicant's involvement in parties or organizations (c) Fraction of asylum seekers for which involvement in military was detected (or not) in their case text (d) Acceptance rates depending on applicant's involvement in military.



FIGURE 5. (a) Fraction of asylum seekers per marital status (b) Acceptance rates per applicant's marital status (c) Fraction of application in which divergences were detected (d) Acceptance rates depending on the detection of divergences.

12) PREVIOUS ASYLUM

Similarly, we have extracted information on whether an applicant has applied for asylum in a different country, prior to their arrival in Denmark. We followed a binary classification which reflects the applicant's seeking of asylum elsewhere, but not the outcome of that application. As shown in Fig. 6(a), such information is rarely detected in our dataset.

13) TORTURE CASES

Some cases are tagged as related to torture, or punishment. For these, we constructed regular expressions and applied language processing for extracting whether the applicant's claims for torture were validated with medical examination. As we report in Fig. 6(c), not many of the cases in our dataset report torture. From those that do, most seem to not have the claim supported.

B. VARIATIONS IN RECOGNITION RATES

We present now the variations in the recognition (or overturn) rates detected in our dataset, that is, at what fraction within each category (i.e., feature) the initial decision by the Danish Immigration Service to deny asylum was overturned by the Refugee Appeals Board.

In Figure 1(b) we present the variations with regards the applicants' country of origin. We observe that the applicants from Eritrea are the only group within which rejection and overturn rates are close to equal. Relatively high overturn rates are observed among Syrian and Ethiopian applicants. The absolute rejection rates in cases such as Morocco and

VOLUME 10, 2022

Libya, call for further investigation, as to whether those rates are justified by the generally lower risk profiles of these nationalities, since the number of persons seeking asylum from these countries is not negligible, as shown in Fig. 1(a). Another interesting point is the absolute rejection of applicants that have not disclosed their country of origin.

Figure 1(d) shows the fraction of granted and non-granted applications per identified gender of the asylum seeker. We observe an ever so slight higher rate of overturn cases in the group of female applicants. We remark that this is not necessarily a sign of bias since gender-related persecution may itself be part of the asylum motive and women may be more at risk of being exposed to persecution than men in certain countries.

With regards to the applicant's identified religion, as per the distribution presented in Fig. 2(a), we focus our attention on Muslim religions, Christianity, Yarsan, Non-religious and Other labels, as every other label has a much smaller representation in the dataset. As shown in Fig. 2(b), among the highlighted religions, the lowest overturn rate is found among Other Muslims, whereas specific Muslim religions, such as Shia and Sunni, enjoy similar recognition rates to Non-religious and Christians. Persecution of religious underrepresented groups - as well as persons who renounce or convert to another religion - is an increasingly common phenomenon around the world. Yet a more detailed analysis, correlating religion with e.g. nationality and/or specific geographic origin, would be needed to better understand the underlying reasons for these rates.



FIGURE 6. (a) Fraction of asylum seekers that have (not) previously applied for asylum in a country other than Denmark (b) Acceptance rates depending on whether an applicant had applied for asylum in a different country before (c) Fraction of asylum seekers that report (or not) torture in their application, and whether such was confirmed or not after medical examination (d) Acceptance rates depending on applicant's submission to torture, and validation of such claims.

Next, we looked at the rates per year of entry of an applicant in Denmark. As shown in Fig. 2(d), no clear-cut pattern can be observed. What we observe is an increase in the overturn rates for applicants that entered in the early 2000s and around the year 2012. The year an applicant entered the country and the year their case was tried by the Refugee Appeals Board are not necessarily the same. It is interesting therefore to examine whether specific years when the decisions were made have had higher recognition rates than others. This is presented in Fig. 7. Indeed, a higher number of granted cases at the appeals level appear for the year 2013. Different factors may help explain this shift. For example, the composition of the Refugee Appeals Board changed in 2013, introducing new members with different professional backgrounds as part of the decision making process. Yet, the fact that the overturn rate drops again already the following year, and does not change significantly as the composition of the board changes once again in 2016, does not support this conclusion. A more likely interpretation of this peak is that it signifies a moment of "legal evolution" [45], where the assessment or legal interpretation in regard to specific types of asylum cases changes. This would typically happen at the appeals level, leading the Refugee Appeals Board to overturn and/or reopen a higher number of cases until the new interpretation has been aligned at the level of the Danish Immigration Service.



FIGURE 7. Fraction of recognition rates per date of decision.

With regards to the applicant's ethnicity, we observe relatively high overturn rates among applicants with identified ethnicity as Bayat, Lorr, Ingustia and Tumal, but also among Persians, Somali, Kurds and Chechen people, as shown in Fig. 3(b).

Moving to the applicant's involvement with political parties or/and organizations, we find no interesting observations in Fig. 4(b), rather a slightly higher overturn rate when no such involvement was detected. Similar results are taken with regards to an applicant's involvement with military, as shown in Fig. 4(d).

A slightly higher percentage of overturned cases is found among applicants that are not, or used to be married, as presented in Fig. 5(b). In Figure 5(d), we observe an almost negligible higher overturn rate within cases where divergences were not detected, or were not reported, in the applicant's decision text. Similar, with regards to an applicant's earlier application for asylum to a country other than Denmark (Fig. 6(b)), a slightly higher overturn rate is observed for applicants that either haven't applied elsewhere, or where such information was not detected in the decision text.

Finally, when looking at cases where torture was reported by the applicant, we find that twice as many applicants were granted asylum when their claim was supported, as opposed to those whose claim was not (Figure 6(d)). Even though the population of torture cases in our dataset is not enormous (see Fig. 6(c)), the observed variation in the torture cases can clearly indicate a (reasonable) favoritism to asylum seekers that have been victims of torture.

Discussion on the Found Variations: Although such variations (where observed) could be the result of biases in the decision process, such conclusion is neither easy, nor accurate to draw due to a number of reasons. Firstly, as already highlighted, our dataset is not fully representative. It does not include cases granted asylum at the first instance. Moreover, not all decisions by the Refugee Appeals Board are made publicly available in the FLN database. Thirdly, and perhaps more importantly, our dataset only includes features directly tied to an asylum seeker's application, which, to some extent, are *legally relevant*. In other words, factors such as the applicant's religion, gender, nationality and year of entry in a country may all have a direct impact on whether the person is at risk of persecution and meet the legal thresholds established by international refugee and human rights law. For instance, the 1951 Refugee Convention specifically requires that a person has a well-founded risk of persecution "for reasons of race, religion, nationality, membership of a particular social group or political opinion". Our analysis thus differs from related works focusing on factors that, from a legal perspective, ought to be irrelevant to the decision of an asylum case, such as the judge's identity, external events or case processing number. This is not to say that bias may not be linked to core applicant characteristics such as nationality and religion [46], but such a conclusion would require a more in-depth qualitative analysis of the cases and/or access to comparable datasets from other countries. We are, at the moment, in the process of acquiring and analysing a more complete dataset of FLN decisions and comparable datasets from around the globe.

We conduct next a more in-depth analysis and understanding on the existence of predictable information in the asylum decision summaries.

IV. MACHINE LEARNING FOR PREDICTING THE OUTCOME OF A CASE

We study now the predictability of the decision outcomes, on the basis of the features considered in our work. First, we calculate the importance of each feature for the predictability of the outcome, using a random forest classifier. Next, we apply a number of ML classifiers on the dataset, giving as predictors every possible combination of our feature space.

Due to the categorical type of most of our features, we constructed a numerical version of each feature in the dataset. That is, binary features, such as the decision, gender, inclusion in political or military organizations, have been coded using 0 and 1 values. Features with a larger range of values, such as the applicant's country of origin, religion and ethnicity, were coded using numbers in the range of the feature's size.

As we observed in Fig. 3(c), our dataset is skewed, with regards to the decision outcomes. Such imbalance could make it difficult for any significant patterns to be revealed. We therefore created a second, balanced dataset, where an equal number of cases of either outcome were sampled from the complete dataset. We discuss our findings on both datasets, referring to them as complete and sampled dataset.

A. FEATURE IMPORTANCE

We use a Random Forest Classifier to calculate the importance score that each feature has in predicting the outcome of an asylum case. Figure 8 shows the importance scores, calculated using the library Scikit-learn in python on the complete dataset. We observe that Religion, the Asylum Category and the Date of Entry of the applicant in Denmark are the top-3 predictors, with the applicant's Ethnicity, Nationality (Country of Origin) and Date of the case's Decision following next. Similar results were achieved on the sampled dataset, with the applicant's Nationality ranking as the 3rd highest feature and ethnicity as the 5th. This analysis suggests that



FIGURE 8. Feature importance scores for predicting the outcome of an asylum case.

some of these features could be carrying information that can be used for predicting a case's outcome.

B. CLASSIFICATION MODELS

Next, we build classifiers to investigate whether we could predict the outcome of an asylum case, and with what accuracy, by looking at some of the applicant's characteristics. We repeat that this predictability may or may not imply a bias in the system, as some of the considered features are legally relevant for the outcome of a case.

We build a number of classifiers, each of them using as input a different combination of features, and a different learning method. Given our 13 features, we ended up with a total of $2^{13} \approx 8000$ combinations. For testing each classifier, we randomly divide our dataset in a training subset (containing 80% of the cases) and a validation subset (containing the remaining 20% of the cases) and applied 10-fold cross validation.

We use the following learning methods [47]: Decision Tree, Random Forest, Support Vector Machine, Logistic Regression, Neural Networks and Naive Bayes. We use the algorithms implemented with the python library sklearn. For Neural Networks, we use Keras.⁴

An overwhelming large number of models and learning methods achieved accuracy of at least 70%, suggesting that certain features (or combinations of features) can predict the cases' outcomes decided by the Refugee Appeals Board. Such predictability constitutes a worrisome finding, with controversial implications for the trust of asylum applicants in the relevant systems. We look next at the results in detail.

1) NAIVE BAYES

All Naive Bayes models on the sampled dataset achieved accuracy between 44% and 62%. The Asylum Type, Date of Decision, Entry date, applicant's gender and involvement

⁴https://keras.io/

in political parties are the features that jointly achieved the highest accuracy. On the complete dataset the accuracy was around 78%, with negligible variance. By looking closely at the performance results, in the complete dataset we notice significantly higher accuracy and precision on the class "Not granted asylum" whereas the class "Granted asylum" had lesser performance. On the sampled dataset, the accuracy and precision values within the two classes where similar, and of at least 46% precision. This suggests that the models primarily guess the outcome and do not really use any input feature as informative predictor.

2) DECISION TREE

Accuracy between 44 - 78% is achieved by the models built using a Decision Tree on the sampled dataset and between 83% and 92% in the complete dataset. Interestingly, on the sampled dataset, most models achieve accuracy on the highest scores (at least 65%) where the accuracy and the precision of both classes is at least 64%. The within-class accuracy and precision is lower in the case of the complete dataset, with values of 45%. The models achieving the higher accuracy on both datasets are trained on combinations of most features we consider in the study, but most of these models contain the applicant's Nationality, ethnicity, religion and date of entry in Denmark. Other features that appear in the higher ranking models are the Gender of the applicant and their involvement in military and in political parties.

RANDOM FOREST

Similar results as the Decision Tree where achieved with a Random Forest classifier, but with the highest achieved accuracy at 82% in the sampled dataset, and 85% in the complete dataset. The same subsets of features as for the Decision Tree are the ones achieving the best performance. The features achieving the highest performance are indeed the ones scoring highest in Fig 8. Although the performance values cannot be used to concretely establish bias in the decision system, these are nevertheless worth highlighting for further exploration, as an at least 78% precision is achieved with those.

4) SUPPORT VECTOR MACHINE

Accuracy between 42% and 61% is achieved when we use SVMs on the sampled dataset. On the complete dataset, however, the accuracy varies from 22% to 78%. By looking closely at those results, we observe that in the sampled dataset, the dominant class "Not-granted asylum" exhibits poor accuracy and precision as opposed to the secondary class, whereas the exact opposite results are observed in the complete dataset. Given the poor in-class performance of the SVMs models, we cannot draw any conclusions and further investigation of the model is necessary.

5) LOGISTIC REGRESSION

Logistic Regression models returned results comparable with Naive Bayes on both datasets.

6) NEURAL NETWORK

Using the library Keras, we built a neural network for predicting the outcome of an asylum case, given the considered features. We tried a number of configurations, in terms of number of hidden layers and length of each layer. Accuracy of at most 69% was achieved by the models that had 2 or 3 hidden layers and 95-120 nodes each, on the sampled dataset. On the complete dataset, the same configuration achieved an accuracy of at most 77%. For the hidden layers, we used the activation function relu, whereas for the output layer we used a sigmoid activation function. Our networks ran in 100 epochs, of batch size 30. These results motivate us for further future exploration and deeper analysis of the neural network approach.

An overview of the classification results on the sampled dataset can be seen in Fig. 9. We omit the confusion matrices, due to the large number of models employed and the variability of the matrices contents across the models, which prevent us from aggregating them in a meaningful way.



FIGURE 9. Boxplots of classification accuracy over all built models, across all tested classifiers, and on the sampled dataset.

V. DISCUSSION

A number of observations can be made from our analysis, in the context of using Machine Learning for predicting outcomes in the legal domain, or/and building models for automated decision-making, in the same domain.

- First, we underscore the need for representative datasets.
- Secondly, we highlight the value of using incomplete datasets, like the one employed in our study, for revealing and studying variations in adjudications.
- Third, we remark on how the choice of the classifier significantly affects the prediction accuracy. Although this does not constitute news in the ML community, it should be considered and discussed when presenting classification and prediction results, especially when the end goal is to conclude to discrepancies caused by biases in the system. Unfortunately, a lot of relevant studies in literature opt for reporting the best results from one classifier without discussing the use of alternate models. In our work, we employed the most popular prediction classifiers, with Random Forest and Decision Tree being the algorithms used in the majority of similar researches

(for instance, references [16], [17]). We found that the large variation of the accuracy results across algorithms, makes the choice of the algorithm an important decision to discuss.

- Fourth, our analysis suggests that some features carry information with potentially predictive properties with regards to the outcome of an asylum case. Information such as an applicant's cultural background, the year they entered the country and/or the year their case was appealed, could be used to predict whether their initially rejected case would be overturned by the Refugee Appeals Board, with as high as 82% accuracy, when a Random Forest Classifier is used. However, such features are to some extent legally relevant for the decision on granting asylum to an applicant. This fourth observation calls for further analysis, in order to validate the predictability of the outcome on the basis of such (or other) features, on a more representative dataset, which is left for future work.
- Last, we would like to define the borders of our analysis. Even though there are international and European treaties which establish the thresholds within which granting or refusing asylum should be practised, individual countries decide on how to assess asylum claims, resulting in possible variability on recognition rates for seemingly similar applications. What is more, asylum decision making entities vary from country to country, from single judges to boards of also varying capacity. Consequently, the results of our study, as well as studies similar to ours, cannot be directly applied or used to explain the recognition rates of other countries. We hope, however, that our work and works similar to ours will inspire more researches across the globe to explore the recognition rates of asylum cases in their countries, resulting in a large database of international asylum law which could serve for in depth analysis of the topic.

Next, we raise some additional points, with regards to asylum case law.

The widespread variations in recognition rates for asylum claims constitute a fundamental puzzle in refugee research: Why, despite decades of regional harmonization and growing international jurisprudence, are national adjudicatory outcomes not more aligned [48], [49]? While political shifts have led to a host of changes in national immigration laws, leg-islative changes rarely impact the structure of procedures or criteria for awarding asylum [50]. Asylum decision-making, or refugee status determination, is moreover a complex process revolving around not only law, but also inter-subjective assessments of applicants' credibility and the import of medical and other forms of evidence. Consequently, the asylum process often appears as "black boxed" to both applicants and scholars [51] and little is known about how different aspects interact and shape outcomes.

Provided access to suitable datasets of national asylum case law, the use of ML methods represents a unique inroad to

interrogate this process and the legal outcomes emanating from it. In areas such as asylum law, where decision making remains non-transparent, subject to administrative discretion and/or more inter-subjective assessments of how the law should be applied, ML may help both researchers/social organisations who wish to identify systemic issues and biases across large numbers of legal decisions as well as courts and legal institutions who seek to preempt bias arising in their institutional practices [52].

VI. CONCLUSIVE REMARKS

Asylum decision-making, or refugee status determination, is a complex process. The use of ML appears to be an efficient way to interrogate this process and the legal outcomes that result from it, granted of course the availability of representative datasets. In this study we analysed a publicly available dataset containing a large number of summaries of asylum cases, initially rejected, and re-tried by the Refugee Appeals Board in Denmark. We highlight variations in the recognition rates, with regards to a number of applicants' features, and apply ML classification in order to study the predictability of the cases' outcomes. We conclude that the choice of applied classifier shapes the predictability outcome. Being in the process of acquiring and analysing an even larger and representative dataset of asylum cases treated in Denmark and in other countries, we are on our way to validate and compare the present results with new findings. Given the more detailed dataset, we plan to identify and extract more elaborate features, and study the predictability of the decision on the basis of those.

REFERENCES

- [1] E. Ash, D. Chen, C. Andrus, D. Godevais, and G. Ng, "Machine prediction of appeal success in U.S. Asylum courts," Tech. Rep. [Online]. Available: https://www.nber.org/~dlchen/papers/Machine_Prediction_of_Appeal_ Success_in_US_Asylum_Courts.pdf
- [2] D. L. Chen, T. J. Moskowitz, and K. Shue, "Decision making under the gambler's fallacy: Evidence from asylum judges, loan officers, and baseball umpires," *Quart. J. Econ.*, vol. 131, pp. 1181–1242, Mar. 2016.
- [3] C. Coglianese and D. Lehr, "Regulating by robot: Administrative decision making in the machine-learning era," *Geo. LJ*, vol. 105, p. 1147, Feb. 2016.
- [4] O. Eren and N. Mocan, "Emotional judges and unlucky juveniles," Amer. Econ. J., Appl. Econ., vol. 10, no. 3, pp. 171–205, Jul. 2018.
- [5] I. Giuffrida, "Liability for AI decision-making: Some legal and ethical considerations," *Fordham L. Rev.*, vol. 88, p. 439, Sep. 2019.
- [6] D. M. Katz, M. J. Bommarito, II, and J. Blackman, "Predicting the behavior of the supreme court of the united states: A general approach," 2014, *arXiv*:1407.6333.
- [7] M. Veale, M. Van Kleek, and R. Binns, "Fairness and accountability design needs for algorithmic support in high-stakes public sector decisionmaking," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, Apr. 2018, pp. 1–14.
- [8] M. Kuziemski and G. Misuraca, "AI governance in the public sector: Three tales from the frontiers of automated decision-making in democratic settings," *Telecommun. Policy*, vol. 44, no. 6, Jul. 2020, Art. no. 101976.
- [9] J. C. Weyerer and P. F. Langer, "Garbage in, garbage out: The vicious cycle of AI-based discrimination in the public sector," in *Proc. 20th Annu. Int. Conf. Digit. Government Res.*, Jun. 2019, pp. 509–511.
- [10] L. T. Rose and K. W. Fischer, "Garbage in, garbage out: Having useful data is everything," *Meas., Interdiscipl. Res. Perspective*, vol. 9, no. 4, pp. 222–226, 2011.
- [11] N. H. Møller, I. Shklovski, and T. T. Hildebrandt, "Shifting concepts of value: Designing algorithmic decision-support systems for public services," in *Proc. 11th Nordic Conf. Hum.-Comput. Interact., Shaping Experiences, Shaping Soc.*, Oct. 2020, pp. 1–12.

- [12] A. Heyes and S. Saberian, "Temperature and decisions: Evidence from 207,000 court cases," *Amer. Econ. Journal: Appl. Econ.*, vol. 11, pp. 65–238, Apr. 2019.
- [13] J. Reis, P. E. Santo, and N. Melão, "Impacts of artificial intelligence on public administration: A systematic literature review," in *Proc. 14th Iberian Conf. Inf. Syst. Technol. (CISTI)*, Jun. 2019, pp. 1–7.
- [14] D. Lehr and P. Ohm, "Playing with the data: What legal scholars should learn about machine learning," UCDL Rev., vol. 51, p. 653, Jun. 2017.
- [15] H. Surden, "Artificial intelligence and law: An overview," *Georgia State University Law Review*, vol. 35, p. 1305, Jun. 2018.
- [16] D. L. Chen and J. Eagel, "Can machine learning help predict the outcome of asylum adjudications?" in *Proc. 16th, Ed., Int. Conf. Articial Intell. Law*, Jun. 2017, pp. 237–240.
- [17] M. Dunn, L. Sagun, H. Şirin, and D. Chen, "Early predictability of asylum court decisions," in *Proc. 16th, Ed., Int. Conf. Articial Intell. Law*, Jun. 2017, pp. 233–236.
- [18] J. Ramji-Nogales, A. I. Schoenholtz, and P. G. Schrag, "Refugee roulette: Disparities in asylum adjudication," *Stan. L. Rev.*, vol. 60, p. 295, Jun. 2007.
- [19] S. Rehaag, "Patrolling the borders of sexual orientation: Bisexual refugee claims in Canada," *McGill LJ*, vol. 53, p. 59, Jun. 2008.
- [20] A. J. Rottman, C. J. Fariss, and S. C. Poe, "The path to asylum in the U.S. and the determinants for who gets in and why," *Int. Migration Rev.*, vol. 43, no. 1, pp. 3–34, Mar. 2009.
- [21] F. E. Marouf, "Implicit bias and immigration courts," New Eng. L. Rev., vol. 45, p. 417, Dec. 2010.
- [22] S. Rehaag, "The role of counsel in Canada's refugee determinations system: An empirical assessment," *Osgoode Hall LJ*, vol. 49, p. 71, Sep. 2011.
- [23] M. Oswald, "Algorithm-assisted decision-making in the public sector: Framing the issues using administrative law rules governing discretionary power," *Phil. Trans. Roy. Soc. A, Math., Phys. Eng. Sci.*, vol. 376, no. 2128, Sep. 2018, Art. no. 20170359.
- [24] M. Suksi, "Administrative due process when using automated decisionmaking in public administration: Some notes from a Finnish perspective," *Artif. Intell. Law*, vol. 29, pp. 87–110, Mar. 2021.
- [25] A. Brown, A. Chouldechova, E. Putnam-Hornstein, A. Tobin, and R. Vaithianathan, "Toward algorithmic accountability in public services: A qualitative study of affected community perspectives on algorithmic decision-making in child welfare services," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2019, pp. 1–12.
- [26] L. McGregor, D. Murray, and V. Ng, "International human rights law as a framework for algorithmic accountability," *Int. Comparative Law Quart.*, vol. 68, p. 309, Apr. 2019.
- [27] L. Ashander, J. Kliestikova, P. Durana, and J. Vrbka, "The decisionmaking logic of big data algorithmic analytics," *Contemp. Readings Law Social Justice*, vol. 11, no. 1, pp. 57–62, 2019.
- [28] S. Delacroix, "Computer systems fit for the legal profession?" Legal Ethics, vol. 21, no. 2, pp. 119–135, Jul. 2018.
- [29] N. H. Møller, I. Shklovski, and T. T. Hildebrandt, "Shifting concepts of value: Designing algorithmic decision-support systems for public services," in *Proc. 11th Nordic Conf. Hum.-Comput. Interact., Shaping Experiences, Shaping Soc.*, New York, NY, USA, Oct. 2020, pp. 1–12.
- [30] B. Lepri, N. Oliver, E. Letouzé, A. Pentland, and P. Vinck, "Fair, transparent, and accountable algorithmic decision-making processes," *Philosophy Technol.*, vol. 31, no. 4, pp. 611–627, Dec. 2018.
- [31] P. Molnar and L. Gill, "Bots at the gate: A human rights analysis of automated decision-making in Canada's immigration and refugee system," Int. Hum. Rights Program (IHRP), Citizen Lab, Munk School Global Affairs, Public Policy, Fac. Law, IT3 Lab, Univ. Toronto, Tech. Rep., Sep. 2018.
- [32] D. L. Chen, "Machine learning and the rule of law," Revista Forumul Judecatorilor, 2019, vol. 19.
- [33] J. Pääkkönen, M. Nelimarkka, J. Haapoja, and A. Lampinen, "Bureaucracy as a lens for analyzing and designing algorithmic systems," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, Apr. 2020, pp. 1–14.
- [34] R. Copus, R. Hübert, and H. Laqueur, "Big data, machine learning, and the credibility revolution in empirical legal studies," *See Livermore Rockmore*, vol. 2019, pp. 21–57, May 2019.
- [35] R. Bevan, A. Torrisi, K. Atkinson, D. Bollegala, and F. Coenen, "Efficient and effective case reject-accept filtering: A study using machine learning," in *Proc. JURIX*, 2018, pp. 171–175.
- [36] N. Aletras, D. Tsarapatsanis, D. Preoţiuc-Pietro, and V. Lampos, "Predicting judicial decisions of the European court of human rights: A natural language processing perspective," *PeerJ Comput. Sci.*, vol. 2, p. e93, Oct. 2016.

- [37] M. Medvedeva, M. Vols, and M. Wieling, "Using machine learning to predict decisions of the European court of human rights," *Artif. Intell. Law*, vol. 28, no. 2, pp. 237–266, Jun. 2020.
- [38] M. Možina, J. Žabkar, T. Bench-Capon, and I. Bratko, "Argument based machine learning applied to law," *Artif. Intell. Law*, vol. 13, no. 1, pp. 53–73, Mar. 2005.
- [39] R. Nanda, G. Siragusa, L. Di Caro, G. Boella, L. Grossio, M. Gerbaudo, and F. Costamagna, "Unsupervised and supervised text similarity systems for automated identification of national implementing measures of European directives," *Artif. Intell. Law*, vol. 27, no. 2, pp. 199–225, Jun. 2019.
- [40] L. Smith-Khan, "Why refugee visa credibility assessments lack credibility: A critical discourse analysis," *Griffith Law Rev.*, vol. 28, pp. 1–25, Oct. 2019.
- [41] R. Pakzad, "Opportunities and challenges of emerging technologies for the refugee system," Centre Int. Innov., Emerg. Technol., Refugees Category, World Refugee Council Res. Paper Ser., Tech. Rep., 2019.
- [42] U. Sadl and H. P. Olsen, "Can quantitative methods complement doctrinal legal studies: Using citation network and corpus linguistic analysis to understand international courts," *Leiden J. Int. Law*, vol. 30, p. 327, Jun. 2017.
- [43] H. P. Olsen and A. Küçüksu, "Finding hidden patterns in ECtHR's case law: On how citation network analysis can improve our knowledge of ECtHR's article 14 practice," *Int. J. Discrimination Law*, vol. 17, no. 1, pp. 4–22, Mar. 2017.
- [44] N. Ridi, "The shape and structure of the 'Usable Past': An empirical analysis of the use of precedent in international adjudication," J. Int. Dispute Settlement, vol. 10, no. 2, pp. 200–247, Jun. 2019.
- [45] T. Gammeltoft-Hansen, "Evolutions in international refugee law," Inaugural Professorial Lect., Univ. Copenhagen, Copenhagen, Denmark, Tech. Rep., Apr. 12, 2019.
- [46] E. Montgomery and A. Foldspang, "Predictors of the authorities' decision to grant asylum in Denmark," *J. Refugee Stud.*, vol. 18, pp. 66–454, Dec. 2005.
- [47] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, no. 10, pp. 2825–2830, 2011.
- [48] F. Trauner, "Asylum policy: the EU's 'crises' and the looming policy regime failure," J. Eur. Integr., vol. 38, no. 3, pp. 311–325, 2016.
- [49] D. E. Anker and M. D. Müller, "Explaining credibility assessment in the asylum procedure," *Int. J. Refugee Law*, vol. 19, no. 3, pp. 599–602, Oct. 2007, doi: 10.1093/ijrl/eem051.
- [50] T. Gammeltoft-Hansen, "Refugee policy as 'negative nation branding': The case of Denmark and the nordics," *Danish Foreign Policy Yearbook*, vol. 2017, pp. 99–125, Jan. 2017.
- [51] B. Latour, Pandora's Hope: Essays on the Reality of Science Studies. Cambridge, MA, USA: Harvard Univ. Press, 1999.
- [52] M. Langford, "Taming the digital leviathan: Automated decision-making and international human rights," *AJIL Unbound*, vol. 114, pp. 141–146, Feb. 2020.



PANAGIOTA KATSIKOULI received the Diploma and M.S. degrees in computer engineering and informatics from the Polytechnic University of Patras, Greece, in 2011 and 2013, respectively, and the Ph.D. degree in informatics from The University of Edinburgh, U.K., in 2017. Since 2017, she has been a Researcher in various institutes, such as Inria, University College of Dublin, and Technical University of Denmark. She is currently a Researcher at the University of Copenhagen. Her

research interests include distributed algorithms, blockchain technology, human mobility, data analytics, applications of machine learning, and distributed algorithms for mobility data.

IEEEAccess



WILLIAM H. BYRNE received the Ph.D. degree from the iCourts, University of Copenhagen. He is currently a Postdoctoral Researcher with the iCourts, University of Copenhagen. He is also involved with the research projects data science for asylum legal landscaping (DATA4ALL), algorithmic fairness for asylum seekers and refugees (AFAR), and advancing data science in migration law (Nordasil.). His research interests include international legal theory, sociology of interna-

tional law, public international law, international human rights law, international refugee law, and legal philosophy.



THOMAS GAMMELTOFT-HANSEN received the M.A. degree in political science from the University of Copenhagen, the M.Sc. degree in refugee studies from the University of Oxford, and the Ph.D. degree in international law from Aarhus University. He is currently a Professor of migration and refugee law at the University of Copenhagen and the Director of the Nordic Asylum Law & Data Laboratory, under which he is PI on a number of collaborative projects combining computational

methods and legal analysis in the migration domain. His research interests include interdisciplinary approaches to Nordic and international refugee and migration law.



ANNA HØJBERG HØGENHAUG is currently pursuing the Ph.D. degree with the University of Copenhagen. Her Ph.D. project forms part of the Data Science for Asylum Legal Landscaping (DATA4ALL) Project. In her project, she examines to what extent national interpretations and applications of international law may help explain outcome variations across the Nordic asylum systems. She has a legal background with a specialization in refugee and human rights law. Her research inter-

ests include international refugee law, human rights law, and data science.







NAJA HOLTEN MØLLER is currently an Assistant Professor with the Software, Data, People & Society (SDPS) Section, Department of Computer Science, University of Copenhagen (DIKU). Her work is interdisciplinary: she is interested in how we can improve our understanding of cooperative work in complex, professional work domains, and the relationship with technology-support.

TRINE RASK NIELSEN is currently pursuing the Ph.D. degree with the Software, Data, People and Society Research Section (SDPS), Department of Computer Science (DIKU), University of Copenhagen, and a part of the Confronting Data Co-Laboratory. Her Ph.D. research is grounded in computer-supported cooperative work (CSCW) and critical data studies.

HENRIK PALMER OLSEN is currently a Professor with the iCourts, University of Copenhagen. His research interests include jurisprudence, human rights, international courts, and data science applied to legal analysis.



TIJS SLAATS received the M.Sc. degree in information technology and the Ph.D. degree in computer science from the IT University of Copenhagen, under supervision of Thomas Hildebrandt. For his Ph.D., he worked on the Technologies for Flexible Cross-Organizational Case Management Systems (FLExCMS) industrial Ph.D. project, in close collaboration with Exformatics A/S, a Danish provider of electronic case management systems. He is currently an Assistant Professor at

the Human-Centered Computing Section, Department of Computer Science, Copenhagen University. His main research interests include business process management, with a particular focus on declarative and hybrid process notations.