






Graphemic Normalization of the Perso-Arabic Script

Raiomond Doctor, Alexander Gutkin,
Cibu Johny, Brian Roark & Richard Sproat

Abstract. Since its original appearance in 1991, the Perso-Arabic script representation in Unicode has grown from 169 to over 440 atomic isolated characters spread over several code pages representing standard letters, various diacritics and punctuation for the original Arabic and numerous other regional orthographic traditions (Unicode Consortium, 2021). This paper documents the challenges that Perso-Arabic presents beyond the best-documented languages, such as Arabic and Persian, building on earlier work by the expert community (ICANN, 2011; 2015). We particularly focus on the situation in natural language processing (NLP), which is affected by multiple, often neglected, issues such as the use of visually ambiguous yet canonically nonequivalent letters and the mixing of letters from different orthographies. Among the contributing conflating factors are the lack of input methods, the instability of modern orthographies (e.g., Aazim, Mansour, and Pournader, 2009; Iyengar, 2018), insufficient literacy, and loss or lack of orthographic tradition (Jahani and Korn, 2013; Liljegren, 2018). We evaluate the effects of script normalization on eight languages from diverse language families in the Perso-Arabic script diaspora on machine translation and statistical language modeling tasks. Our results indicate statistically significant improvements in performance in most conditions for all the languages considered when normalization is applied. We argue that better understanding and representation of Perso-Arabic script variation within regional orthographic traditions, where those are present, is crucial for further progress of modern computational NLP techniques (Conneau et al., 2020; Muller, Anastasopoulos, Sagot, and Seddah, 2021; Ponti et al., 2019) especially for languages with a paucity of resources.

1. Introduction

The Modern Perso-Arabic script derives from the fourth century North Arabic script, which in turn was adapted from the Nabatean Aramaic

Raiomond Doctor  0009-0005-8684-4937, on contract from Optimum Solutions, Inc. Alexander Gutkin  0000-0001-6327-4824, Cibu Johny  0009-0008-4220-5414, Brian Roark  0000-0002-7292-6246 & Richard Sproat  0000-0002-9040-5196

Google Research: India, United Kingdom, United States, and Japan
E-mail: {raiomond, agutkin, cibu, roark, rws}@google.com

Y. Haralambous (Ed.), *Grapholinguistics in the 21st Century 2022. Proceedings*
Grapholinguistics and Its Applications (ISSN: 2681-8566, e-ISSN: 2534-5192), Vol. 9.
Fluxus Editions, Brest, 2024, pp. 315–376. <https://doi.org/10.36824/2022-graf-gutk>
ISBN: 978-2-487055-04-9, e-ISBN: 978-2-487055-05-6

script to write the Arabic language (Bauer, 1996; Gruendler, 1993). Due to the spread of Islam throughout much of Africa, Asia, and parts of Europe, it has come, in its various adapted forms, to be one of the most widely used scripts in the modern world. Due to its reasonable flexibility in representing phonological structure, the script was adopted to write a large number of languages spanning diverse language families such as Afro-Asiatic, Indo-European, Niger-Congo, Turkic, and Sino-Tibetan, among others. Adaptations are found as far south as Southern Africa with Arabic having been used for Afrikaans (Kotzé, 2012) and Malagasy (Versteegh, 2001); as far east as East Asia for Chinese (Suutarinen, 2013) and Japanese (Kaye, 1996; Naim, 1971); and into Eastern Europe for writing languages of Muslim Slavs, such as Bosnian (Buljina, 2019). While many of these adaptations have not survived, the Arabic script and its derivatives are still used for scores of languages with a total population of speakers of over 600 million.¹ For some linguistic areas, such as the Dardic languages of Northern Pakistan, most of which were unwritten until very recently, the Perso-Arabic script is the *only* serious contender when developing a new writing system; see for example Torwali (Torwali, 2019),² and Palula (Liljegren, 2016).

While the original Semitic scripts were pure consonantal scripts (*abjads*), three letters—*alif* /ʔ/, *ya* /y/, and *waw* /w/—came to be used as *matres lectionis* to represent long vowels, and further diacritics were developed to (optionally) represent such features as short vowels and gemination (*shadda*), among others (Bauer, 1996).

The original North Arabic script was rather ambiguous, since Arabic had a larger consonant inventory than Aramaic, and some of the consonant letters had to do double duty—a problem exacerbated by the cursivization of the script. The resulting ambiguities were resolved by the use of various numbers of dots over or under the letters to disambiguate the various uses (Bauer, 1996; Kaplony, 2008), a system called *iʿjām* (إِعْجَام). For example the inferior dot in <ب> /b/ distinguishes it from <ن> /n/ with a single dot on top, and then again from <ت> /t/ with two dots on top, and again still from <ث> /θ/ with three on top. Though the set of consonants to be disambiguated is of course limited in Arabic itself, the *iʿjām*, once started, evolved into a productive way to produce new consonant symbols when the script was adapted to new languages. This has consequently allowed languages to have their “own” version of the Perso-Arabic script, where the only difference with the scripts used for a language’s neighbors is in the use of distinctive *iʿjām*-

1. https://en.wikipedia.org/wiki/List_of_writing_systems\#List_of_writing_systems_by_adoption

2. <https://www.blog.google/around-the-globe/google-asia/torwali-language-and-its-new-android-keyboard/>

augmented consonants. This is true, for example, for adaptations of the script to the many Dardic languages, where each has one or two consonant symbols not found in the scripts of its neighbors.

As noted above, the Arabic script and its derivatives include diacritics that allow one to specify all vowels and other phonetic features such as gemination. However in the normal daily use of the script for Arabic and other languages, these are typically omitted. In Arabic this means that the script is still technically an *abjad*, since the written symbols mostly represent consonants. However to varying degrees, the derived scripts have departed from this, and some of them are full alphabets. Thus according to Kaye (1996), the Persian writing system is an *abjad* as are Urdu and Jawi, the old Malay Arabic-based writing system; however the Kurdish and Uyghur writing systems are alphabets. Parallel developments occurred with adaptations of the Hebrew *abjad* so that Yiddish orthography (Aronson, 1996) is an alphabet.

Historically each geographic region posed its own unique sociolinguistic challenges resulting in the emergence of different adaptation strategies and orthographic traditions across South Asia (Qutbuddin, 2007; Wink, 1991), Southeast Asia (Abdullah et al., 2020; Kratz, 2002; Ricci, 2011), and Africa (Mumin, 2014; Ngom and Kurfi, 2017), among other regions (Castilla, 2019; Suutarinen, 2013). This diversity as well as the flexible nature of the script is reflected in a large and growing inventory of Perso-Arabic code points in the Unicode standard (Unicode Consortium, 2021) with accompanying ambiguities associated with representing the script using the digital medium that we briefly outline in §2. We then provide an overview of some of the regional orthographies for eight languages selected from diverse language families in §3. A significant amount of digital representation ambiguities manifest by these orthographies is resolved computationally using finite-state normalization methods for Perso-Arabic, described in §4, that we developed for this purpose.³ We study the effects of normalization of real-world text using statistical and neural techniques, and present our findings in §5.1 and §5.2, respectively. The code and the results accompanying the experiments have been released.⁴

2. Perso-Arabic in the Digital Medium

As was mentioned previously, an important feature which has led to the adoption of the script by different cultures to use it to transcribe their

3. <https://github.com/google-research/nisaba>

4. https://github.com/google-research/google-research/tree/master/perso_arabic_norm




FIGURE 1. Core *rasm* shapes, or *archigraphemes* according to Milo (2002), of Arabic shown in (a), and examples of Arabic letters derived with *iḥjām* demonstrating their disambiguation function in (b), after (Nemeth, 2017)

language, is its very flexibility.⁵ At its core the Arabic script comprises 18 basic shapes often referred to as *rasm* (رسم) or “drawing” (Daniels, 2013; Kurzon, 2013). These can be modified in various ways: apposing one to four dots (*iḥjām*) placed above, below or inside a character (as shown in Figure 1); using modifier signs such as the subscript or superscript *small hamza*; placing diacritics or *tasbkīl* (تشكيل) and in certain cases even adding a new shape based on the basic Arabic template. Thus for example Urdu, discussed in more detail in §3.1, substantially expanded the original Arabic writing system adapting it to its phonology by introducing additional *iḥjām* characters, modifiers, and even creating new shapes such as the *bari yeb* <ع> or the *heb do chashmee* <ه> for handling aspiration.

Similar to Brahmic scripts, the Perso-Arabic script often provides more than one way to compose a character in the digital medium (Unicode Consortium, 2021). For example, the *alef with madda above* letter can be composed in two ways: as a single character <آ> (U+0622) or by adjoining *madda above* to *alef* (U+0627 followed by U+0653). This results in presentation ambiguity and the Unicode standard provides a certain number of canonical normalization forms, such as the Normalization Form C, or NFC, to handle such cases (Whistler, 2021). A normalization process is required to convert strings to such canonical forms. In contrast to Brahmic script normalization, where atomic forms are normalized to their decomposition, Perso-Arabic normalization in NFC involves combining decomposed elements into a single glyph. Thus the individual *alef* and *madda above* will be normalized to a single glyph. Our investigations have found many cases of this kind of *visual ambiguity* in the Perso-Arabic script beyond what is covered in NFC.

Some of these visual ambiguities are illustrated by a simple example in Table 1, where six alternate representations for the Arabic word for “leader” are shown rendered in Naskh style along with the corre-

5. Scripts are sets of characters used jointly in written representation while writing systems additionally consist of the rules and conventions used when employing a script for a particular language.

TABLE 1. Six different spellings of the Arabic word for “president” (MSA: /ra.ʔi:s/) rendered in Naskh. For each row, the Unicode character differences with the Unicode string in the first row are highlighted. The last column indicates the type of transformation required to bring the relevant Unicode string to the canonical form displayed in the first row of the table.

Display	C ₁	Unicode Character Sequence				C ₅	Transformation
رئيس	<i>reb</i>	<i>yeb with hamza above</i>	<i>yeb</i>	<i>seen</i>			
	U+0631	U+0626	U+064A	U+0633			
رئيس	<i>reb</i>	<i>yeb</i>	<i>hamza above</i>	<i>yeb</i>	<i>seen</i>		Unicode NFC
	U+0631	U+064A	U+0654	U+064A	U+0633		
رئيس	<i>reb</i>	<i>alef maksura</i>	<i>hamza above</i>	<i>yeb</i>	<i>seen</i>		Visual Normalization
	U+0631	U+0649	U+0654	U+064A	U+0633		
رئيس	<i>reb</i>	<i>yeb with hamza above</i>	<i>farsi yeb</i>	<i>seen</i>			Visual Normalization
	U+0631	U+0626	U+06CC	U+0633			
رئيس	<i>reb</i>	<i>farsi yeb</i>	<i>hamza above</i>	<i>yeb</i>	<i>seen</i>		Reading Normalization
	U+0631	U+06CC	U+0654	U+064A	U+0633		
رئيس	<i>reb</i>	<i>farsi yeb</i>	<i>hamza above</i>	<i>farsi yeb</i>	<i>seen</i>		Reading Normalization
	U+0631	U+06CC	U+0654	U+06CC	U+0633		

sponding Unicode sequences ranging from four to five characters. The spelling in the first row of the table is the correct spelling of the word in Modern Standard Arabic (MSA) orthography. The visual forms in the second, third and fourth rows are visually identical to the correct spelling, but are represented digitally as distinct sequences of characters. The second example, while ambiguous, is handled by the Unicode NFC normalization, which brings it to the canonical form provided by the first row by rewriting a (decomposed) sequence of *yeb* and *hamza above* as its canonical single-letter counterpart *yeb with hamza above*. The form in the third row is more problematic. It arises from a five-character sequence which has *alef maksura* instead of *yeb* and is also visually identical to words in rows one and two. However, unlike the spelling in the second row, the sequence *alef maksura* followed by *hamza above* does not have a canonical composed form under Unicode.⁶ Hence, while this

6. See the discussion in (Pournader, 2010) on how this came about.

form is visually identical to the first and second forms, it is treated as a distinct word in the digital medium. Similarly, the example in the fourth row is visually identical to the preceding examples, but arises due to using *farsi yeb* instead of *yeb* which is illegal under MSA orthography. No standard transformation is provided by the Unicode to cope with cases like this example as they are orthography- and language-specific. We refer to the class of normalizations that result in something that is visually identical as *visual* normalization (Johny, Wolf-Sonkin, Gutkin, and Roark, 2021).

The second group of ambiguities is illustrated by the last two rows of Table 1. The fifth and sixth examples, while visually identical to each other, differ slightly from the reference in extra *ījām* dots produced by combining *farsi yeb* with *hamza above* word-internally. Similar to the examples in rows three and four, no fallback normalization strategy is provided by the Unicode standard to handle such cases as it is not clear without prior context which orthography is intended. We refer to the class of normalizations that result in something that is not visually identical as *reading* normalization (Gutkin, Johny, Doctor, Wolf-Sonkin, et al., 2022).

As we mentioned above, unlike the canonical Unicode transformations, it is impossible to define most of visual and reading normalizations outside the orthographic context. Consider a sequence consisting of *waw* (U+0648) followed by *damma* (U+064F) whose visual form <ّ> is identical to letter *u* (U+06C7) used in Kazakh, Uzbek and Uyghur Perso-Arabic orthographies among some others (Aqtay, 2020; Haralambous, 2021). Normalization of *waw* and *damma* to its “canonical” form *u* should only be performed for these languages. We introduce the visual and reading script normalization framework more formally in §4.

One could argue that the ambiguities described above are not problematic and are the natural outcome of the specific properties of the script (e.g., its cursive form and the presence of positional variants), the vast number of orthographic adaptations and specifics of its implementation in digital medium. However, as we show in §5, the resolution of such ambiguities on a language-by-language basis positively impacts the quality of computational models of natural language. Furthermore, the visual ambiguities manifest by various Perso-Arabic writing systems represent a growing challenge to cybersecurity. From the standards’ perspective this is being gradually addressed by the Internet Corporation for Assigned Names and Numbers (ICANN) by developing a robust set of standards for representing various Internet entities in Perso-Arabic script, such as domain names, with particular focus on visually confusable character variants (ICANN, 2011). In addition, security implications, such as development of mechanisms for protection against *phishing* and *spoofing* attacks, are actively studied by the relevant

cybersecurity literature on Perso-Arabic (Ahmad and Erdodi, 2021; El-sayed and Shosha, 2018; Hussain et al., 2016).

A detailed analysis of the causes for the various types of Perso-Arabic script ambiguities described above are outside the scope of this work. It is however worthwhile to briefly mention some of them here. One of the causes is the relative complexity and several special properties of the Perso-Arabic script itself. The script has four key orthographic properties that are relevant here: (i) relative visual similarity of the *rasm* forms; (ii) *allography*, i.e., positional variants of letters (initial, medial, final and isolated); (iii) cursivity/ligaturing, and (iv) *non-linearity*, the extensive and sophisticated use of various types of *tashkīl* and *iqjām* (Yassin, Share, and Shalhoub-Awwad, 2020). The combination of all these properties was demonstrated to lead to relatively more involved visual processing of the script (compared to languages that use the Latin script) in the psycholinguistics and neuropsychology literature on reading for MSA (Boudelaa, Perea, and Carreiras, 2020; Eviatar and Ibrahim, 2014; Hermena and Reichle, 2020), but also for the Uyghur Perso-Arabic orthography (Yakup, Abliz, Sereno, and Perea, 2015).

The Perso-Arabic script support in Unicode is ever-growing, which is reflected by the number of recent proposals for new characters to better reflect the existing low-resource orthographies (Chitralli, 2020a,b; Evans and Warren-Rothlin, 2018; Patel, Riley, and MacLean, 2019) as well as to ease the encoding of the Quranic literature (Sh., 2022). The process of updating the standard is usually time-consuming, as demonstrated by the case of Torwali, which took two years from the time of the original proposal (Bashir, Hussain, and Anderson, 2006) to encode the missing letter *bab with small arabic letter tab above* <ځ> (U+0772). This letter completed the full character inventory for this emerging orthography in Unicode, which facilitated further developments of linguistic resources (Uddin and Uddin, 2019). As we found in our experiments, in the absence of the required characters, visually confusable variants or sequences of variants from foreign orthographies are often used by the Unicode-compliant input methods and converters from non-Unicode compliant fonts. Conversely, these methods take time to catch up with the Unicode standard once it introduces the missing features. To this one can add multiple confounding factors involved in the modern evolution of orthographies for hitherto unrecorded languages, which leads to rich orthographic diversity even among neighboring languages. For example, according to Bashir (2015, p. 14), the retroflex voiceless sibilant /ʂ/ present in several languages of northwestern Pakistan is represented differently by the regional writing systems: they all share the same *rasm* shape for letter *seen* <س> modified as: (1) *seen with small arabic letter tab and two dots* <ښ> (U+0770) for Khowar; (2) *seen with extended arabic-indic digit four above* <سٲ> (U+077D) for Burushaski; (3) *seen with four dots above* <سٴ> (U+075C) for Torwali; (4) *seen with two dots vertically above* <سٶ> (U+076D) for

Gowri; and (5) the corresponding two distinct characters for Kalasha and Shina remain unrepresented in Unicode.

3. Perso-Arabic Script Diaspora: Selected Language Summaries

In what follows, we briefly delve into some adaptations of Perso-Arabic script. We limit our discussion to Perso-Arabic orthographies of eight languages, some of which are written in several scripts. The five languages from Indo-European family are Central Kurdish (Sorani), Kashmiri, Punjabi, Sindhi and Urdu. Two further languages, South Azerbaijani and Uyghur, come from the Turkic family. Finally, we provide a brief overview of the Perso-Arabic orthography called Jawi for Malay from the Austronesian family. The concise language-specific letter inventories are provided in §A. Our software covers more orthographies, such as Balochi, Dari, Modern Standard Arabic, Pashto, Persian and Uzbek, yet we felt that our choice of the above eight writing systems is representative of the kinds of normalization challenges one is likely to encounter.

3.1. Urdu

Ethnologue classes Urdu as the tenth most spoken language in the world with over 70 million speakers using Urdu as their first language.⁷ The national language of Pakistan, one of the 22 official languages of India, and a registered dialect in Nepal, Urdu is also spoken and used in 30 odd countries.⁸

The origins of Urdu are debatable and some scholars trace it back to the 6th century CE (Schmidt, 2007) but it was the Muslim invasion of Sindh in 711 which acted as a catalyst. By the time of the Mughal Empire and at the end of the 18th century it was the lingua franca around Delhi and was called *Zaban-e-Urdu*,⁹ the word Urdu derived from the Turkic word *ordu* for “army” (Lelyveld, 1994).¹⁰ The expansion of the Sultanate to the south gave rise to Dakhani Urdu (Mohamed, 1968). Urdu has a close association with Hindi since they share a common Indo-Aryan origin. Whereas Urdu is written in Perso-Arabic, Hindi uses Devanagari. The difference is best seen in the two versions of “The Chess Players”

7. <https://www.ethnologue.com/guides/ethnologue200>

8. <https://www.ethnologue.com/language/urd>

9. Urdu was also called *Hindi*, *Hindustani*, *Deblavi*, and *Lasbkari*. But the term Urdu became most acceptable.

10. <https://www.rekhtadictionary.com/meaning-of-urdu>

story written by Munshi Premchand, who authored the same story in both languages (Davis, 2015).¹¹

TABLE 2. Urdu in Nastaliq (top), Naskh (center) and the corresponding transliteration (bottom). Samples taken from a poem (*nazm*) by Faiz Ahmed Faiz

Nastaliq
تیرا غم ہے تو غم دہر کا جھگڑا کیا ہے تیری صورت سے ہے عالم میں بہاروں کو ثبات
Naskh
تیرا غم ہے تو غم دہر کا جھگڑا کیا ہے تیری صورت سے ہے عالم میں بہاروں کو ثبات
Transcription
terā ḡham hai to ḡham-e-dahr kā jhagḏā kyā hai terī sūrat se hai aalam meñ bahāroñ ko sabāt

Urdu as used in India and Pakistan is written in the Nastaliq style—a writing style developed in Iran from the Naskh style around the 13th century. Easy to write by hand, it posed problems when ported to metal type. Digital typography has to a certain extent solved the problem and text can be seen in the Nastaliq style, however media on the Web prefers Naskh (Parhami, 2020). A sample of two lines of a *nazm* by Faiz Ahmed Faiz in Table 2 demonstrates the differences between the two styles.¹²

The Urdu writing system is an *abjad*, borrowed from Persian which in turn is borrowed from Arabic. Persian added four characters (<پ>, <ز>, <گ>, and <چ>) to the 28 basic characters borrowed from Arabic, bringing the total to 32. Persian further modified the character set by replacing the Arabic characters <ي> and <ك> with <ی> and <ک>, respectively. To these 32 Persian characters, Urdu added: (1) the three letters <ٹ>, <ڈ>, and <ظ> to accommodate retroflexes; (2) <ں> to handle nasalization; (3) the *two-eyed be* <ھ> to accommodate the aspirated forms of 17 or 18 letters; (4) the *yeh baree* <ے> to represent /e/ at the end of the words; and (5) <و> *gol be*, also called *choṭī be*. Since diacritics are a systematic component of Perso-Arabic, this was possible without upsetting the graphic equilibrium of the script (Coulmas, 1999, p.560). The added *high hamza* placed above *farsi yeh*, *yeh baree*, *beb goal* and *waw* is used to create additional values and the *teh marbuta* <ة> marks feminine gender for

11. <https://thewire.in/culture/why-the-perso-arabic-script-remains-crucial-for-urdu>

12. <https://www.rekhta.org/nazms/>

nouns and adjectives. A further 5 characters were added to represent the 10 vowel phonemes, and an additional 5 to 10 diacritics were used when precision was needed.¹³

3.2. Punjabi (Shahmukhi)

TABLE 3. Shahmukhi in Nastaliq (top), Naskh (center) and the corresponding transliteration (bottom). Samples taken from a poem by Baba Farid

Nastaliq
کافی آوو سکھی سہلیو مل مسلت گوئے آپو آپنی گل نوں بھر ہنجھو روئے کھڈے لالچ لگیاں میں عمر گوائی کدے نہ پونی ہتھ لے اک تندڑی پائی
Naskh
کافی آوو سکھی سہلیو مل مسلت گوئے آپو آپنی گل نوں بھر ہنجھو روئے کھڈے لالچ لگیاں میں عمر گوائی کدے نہ پونی ہتھ لے اک تندڑی پائی
Transcription
kaafi—aavo sakhi saheliyo mil maslat goiye aapo aapni gal nuun bhar hanjhū roiye khaDe lālach lagyañ maiñ umar gavā.ī kade na puuni hath lae ik tandaḌī paa.ī

The Shahmukhi writing system is used to record Punjabi in the Perso-Arabic script. According to Ethnologue this language is mainly spoken in Pakistan but also in other countries, especially in Punjab in India, with a total number of Punjabi speakers around 66 million.¹⁴ The language is also known as Jangli, Lahanda, Lahnda, Lahndi, Panjabi, Panjabi Proper and Punjabi. Historically Shahmukhi was used by Sufi poets of the then Punjab region. One of the earliest instances of this writing system is its use by the Sufi poet of Punjab, Baba Farid in the 12th century (Singh and Gaur, 2009). After the partition of India, Shahmukhi became the writing system of choice for writing Punjabi by the Muslim population in Punjab. Hindus and Sikhs in the Indian state of

13. The number of characters is debated, cf. <https://www.dawn.com/news/919270>.

14. <https://www.ethnologue.com/language/pnb>

Punjab adopted the Brahmic Gurmukhi script to write Punjabi, giving rise to Eastern Punjabi (Grewal, 2004). The relationship between Shahmukhi and Gurmukhi closely parallels that between Urdu and Hindi. Shahmukhi is an *abjad* and is written from right to left. It was highly influenced by Persian, but the present day writing system was modified to suit the requirements of the Punjabi language and as in the case of Urdu, a considerable number of characters were added. Like Urdu, Shahmukhi favours Nastaliq, but Naskh is used by digital media on the web. A sample of four verses from a poem by Baba Farid in both styles is shown in Table 3.

Urdu and Shahmukhi share the same character set, except that Shahmukhi admits a few more letters. The number of characters in Shahmukhi, like Urdu, is a matter of debate and some scholars admit four more characters <ٲ>, <ٴ>, <ٲ>, and <ٴ> in addition to the retroflex lateral *lla* <ٲ> and the retroflex nasal *nnā* <ٲ> (Bashir and Connors, 2019, pp.62, 77). Of these, our analysis shows that only <ٲ> and <ٲ> seem to be in use. The <ٲ> character is used to mark end-of-word nasals. Like Urdu, the *two-eyed he* <ٲ> is used to accommodate the aspirated forms of 17 or 18 letters. To these can be added the *high bamza* (U+0674) placed above *farsi yeh*, *yeh baree*, *heh goal* and *waw* to create additional values. Finally, five characters are used to represent the 10 vowel phonemes, and an additional 5 to 10 diacritics are used when precision is needed in cases such as consonant clusters or gemination. Shahmukhi and Urdu are thus mutually intelligible as writing systems.

3.3. Sindhi

Sindhi is an Indo-Aryan language spoken by the inhabitants of Sindh in the western part of the Indian subcontinent. It is one of the official languages of Pakistan and one of the 22 scheduled languages in India. Thanks to the Sindhi Diaspora it is spoken in quite a few countries and as per Ethnologue, has over 33 million speakers around the world.¹⁵ Sindhi is recorded both in Perso-Arabic as well as Devanagari scripts. The traveler Al-Biruni in his *Tarikh-al-Hind* states that Sindhi was written in three scripts: Ardhanagari, Mahajani and Khudabadi (Sachau, 1910). But, the standardization of the Sindhi Perso-Arabic writing system (“arabi Sindhi”) dates back to the 19th century. Prior to that, Sindhi Muslims had made attempts to write the language using Arabic, but the formal character set of Sindhi, as it is known today, goes back to 1853 when it was standardized by the British colonial authorities (Dow, 1976) and a set of 52 letters to accommodate the complexities of the sound system of Sindhi was identified. Sindhi is an *abjad* but unlike Urdu or Shah-

15. <https://www.ethnologue.com/language/snd>

TABLE 4. Sample of Sindhi in Naskh (top), Devanagari (center), and the corresponding Latin transliteration (bottom). Samples taken from a poem by Shah Abdul Latif Bhittai

Naskh
سڄڻ ۽ سائيهه ڪنهن اناسي وسري، حيث تينن کي هوءَ، وطن چن وساري.
Devanagari
सज्जन ऐं साणेहु कं हिं अणासी विसरी, हैफ तनीं खे होइ, वतनु जिनि विसारी.
Transliteration
sajanu ain sanehu kanhin anasi visri, haif tanin khe hoi, vatanu jini visari.

mukhi, Sindhi is only written using the Naskh style. A sample stanza from “Shah Jo Risalo” by Shah Abdul Latif Bhittai (Lajwani and Mirjat, 2021) shown in Table 4.

The addition of digraphs and the *hamza* over *yeh* and *waw*, as well as the diacritics to indicate the short vowels placed above *alef* and *waw*, brings the size of modern Sindhi letter inventory to 64 (Lekhwani and Lekhwani, 2014). For short vowels in particular, the following four letters <ا>, <اِ>, <اُ>, <ؤ>, composed by placing diacritic marks *fatha*, *kasra* and *damma* over *alef*, and the *damma* over *waw*, were added. To accommodate the large number of characters in its repertoire, Sindhi modified the Arabic *rasm* by addition of more *ijām* dots.

Certain features of the character set of Sindhi make for the uniqueness of the writing system. Unlike Urdu or Shahmukhi, the *high hamza* is already accommodated over <ؤ> and <ئ>. Sindhi admits four implosives <ڳ> *ga*, <ڄ> *ja*, <ڏ> *da*, <ڀ> *ba*, and two single letter words <ء> *ain* (“and”), and <م> *men* (“in”). Like Urdu, Sindhi has four letters to indicate /z/: <ض>, <ظ>, <ز>, <ذ>; three letters for /s/: <س>, <ص>, <ت>; and two letters for /h/: <ح> and <ھ>. However, unlike Urdu which uses the *two eyed he* or *be do chashmi* <ھ> to mark the aspirates, Sindhi has individual characters for all the aspirates with the exception of <ڱ> *gha*, <ڄھ> *jba*, and <ڙھ> *rba*. Vowel diacritics are not normally used. However if needed Sindhi has three diacritics used to indicate the short vowels. Additional diacritics are used to mark consonant clusters (*sukun*, U+0652) and gemination (*shadda*, U+0651).

TABLE 5. A sample of Kashmiri proverbs

Nastaliq	Transliteration	Translation
کنو کن بتر لادن	Kanav kin batı ladun	Stuff rice through the ears: to overfeed
کنس کنس بتر لادن	Kanas batı ladun	Stuff the ear with rice: advice wasted on a fool

3.4. Kashmiri

Kashmiri is a language from the Dardic family spoken in the Union Territory of Jammu and Kashmir, Himachal Pradesh and its outlying regions (Koul and Wali, 2015). Ethnologue lists around 7 million Kashmiri speakers in India and other countries.¹⁶ It is a statutory language of provincial identity in Jammu and Kashmir¹⁷ and is one of the 22 scheduled languages in India.¹⁸ According to B. B. Kachru (2016), Kashmiri is the only Dardic language with a literary tradition and for which the written records have survived. Kashmiri is one of the three scheduled languages of India that is written using a Perso-Arabic script, the other two being Urdu and Sindhi. As in the case of Sindhi, Kashmiri is also written in Devanagari script, suitably modified to accommodate the sounds of the language.¹⁹

Historically Kashmiri was written in the Sharada script, an *abugida* from the Brahmic family (Khaw, 2015). Sharada fell into disuse because it could not represent the complex sound system of the language. Successive invasions of the region slowly led to the adoption of the Arabic script. By the 14th century, Muslim rule in Kashmir was established and Kashmiri in Perso-Arabic script was adopted (Yatoo, 2012). The writing system evolved with time and the Arabic *rasm* were suitably adapted to add new characters to the repertoire. Today the Perso-Arabic script is recognized as the official writing system for the language. It is written in both Naskh and Nastaliq; and although the latter is favoured as the desired style, digital media prefers Naskh owing to the non-availability of a Nastaliq font for the script. Kashmiri is renowned for its proverbs (R. Kachru, 2021; Koul, 2006) and a sample of two proverbs²⁰ in Nastaliq, with the corresponding transliterations and English translations, is shown in Table 5.

16. <https://www.ethnologue.com/language/kas>

17. In 2020, the Parliament of India passed a bill to make Kashmiri an official language of Jammu and Kashmir along with Dogri, Hindi, Urdu and English.

18. <https://rajbhasha.gov.in/en/languages-included-eighth-schedule-indian-constitution>

19. कौशुर (Koshur)

20. <https://kashmiridictionary.org/kanas-bati-ladun/>

Kashmiri is an *abjad*, but because all vowel sounds are regularly indicated in its orthography, the writing system is somewhat closer to an alphabet similar to Sorani Kurdish and Uyghur. The consonant inventory of Kashmiri consists of 37 letters. Some of these letters are shared with Urdu, Khowar and Shahmukhi orthographies, like the letter *rreh* <ڙ> for representing voiced retroflex flap /ɽ/ or the *ddal* <ڛ> for the voiced retroflex plosive /ɖ/. Kashmiri, like Urdu and Shahmukhi, uses the *two-eyed he* <ه> in the construction of aspirated consonants. However, unlike Urdu and Shahmukhi only six such digraphs are permitted: <هت>, <هڄ>, <هپ>, <هچ>, <هڙ>, and <هڪ>. The character <ي>, a *yeh with a ring below* needs special mentioning. Kashmiri uses <ي> to mark palatalization which is a common feature in the language.

Kashmiri has one of the largest inventories of vowel letters, which are arranged in eight pairs of short and long vowels.²¹ Kashmiri uses the *kasra*, *damma* and *fatha* as short vowel diacritic markers. Kashmiri modifies the *waw* to add new vocalic values: <و> *waw with ring* to represent the sound /ɔ/; <ۆ> *waw with inverted v on top* for /o/; and <ؤ> *waw with inverted damma* for a long /u:/. Additionally, *yeh baree* with an *inverted small v* marks the short /e/. Two combining marks are unique to Kashmiri, these are the *wavy hamza above* and *wavy hamza below*. The first is always used in conjunction with *alef* and represents a long schwa /ə:/, while the second is used along with *alef* to represent /i:/.

Similar to Urdu and Shahmukhi, Kashmiri nasalisation is marked by the *noon ghunna* <ٲ> which can only occur in final position. When in medial position it is replaced by the letter *noon* <ن>. Kashmiri also uses two other combining marks to mark gemination using *shadda*, and *sukun/jazm* (also called a *vowel killer*) to mark consonant clusters. The rendering of standard *sukun* diacritic (U+0652) is unique to Kashmiri writing system and has the shape of an inverted <v>.

3.5. Central Kurdish (Sorani)

Sorani is the Perso-Arabic writing system used to write the Kurdish language in Iraq, mainly in Iraqi Kurdistan (Haig, 2018). This Indo-Iranian language is also spoken in regions adjoining Iran and Turkey. Ethnologue identifies three geo-linguistic variants of the language depending on where it is used: Central (Zimane Sorani), Northern (Kurmancî) and Southern (Kurdi Xwarîn or Pehlewani). The name Sorani derives from the Soran Emirate, located in the area known today as Iraqi Kurdistan. Ethnologue lists around 4.7 million Sorani speakers in Iraq and total

21. According to <https://r12a.github.io/scripts/arabic/ks.html>. Slightly different inventory is provided in <https://kashmiridictionary.org/category/learn-kashmiri/vowels-learn-kashmiri/>.

TABLE 6. Sorani sentence in Naskh and Latin translating as “Twenty million people are asking for tickets to participate in the last Led Zeppelin concert”

Naskh
<p>بیست ملیۆن کەس داواى بلیت دەکەن بۆ بەشداری کردن له دوا کۆنسێرتی لێد زێپالین</p>
Latin
<p>bîst miliyon kes dawayi bilît deken bo beşdarî kardîn lah dawa konisêrtî lêd zêpalîn</p>

number of Sorani speakers in all countries at 5.3 million.²² Unlike other languages in this study, Kurdish is not recognized today as the official language in any of the regions where it is used,²³ despite popular movements by Iraqi Kurds to give the language an official status.

Sorani in Perso-Arabic traces its origins to the Sulaimani region. The first trace of this writing system is found in “Mahdîname” by Mullah Muhammad Ibn ul-Haj completed circa 1762 (Bozarslan, Gunes, and Yadirgi, 2021). The rise of the Baban dynasty encouraged the growth of Sorani and it became a medium for prose and poetry (Khalid, 2015). This continued until the Baban dynasty was overthrown around 1856. However, under British rule in the 19th century, Sorani literature and journalism flourished and multiple attempts were made to standardize the writing system, which led to eventual codification of the Sorani alphabet in the 1920s (Campbell, 1994).

Unlike most other Perso-Arabic writing systems, Sorani is a true alphabet, the vowels being explicitly marked. Sorani is written in Naskh style. A sample text from the Pewan corpus (Esmaili et al., 2013) is provided in Table 6.²⁴ As is the case with all languages adopting Perso-Arabic, in order to represent the phonemic features of the language, Sorani has evolved a system of letters some of which are unique to this writing system. These include the three unique consonant letters that are constructed by adding dots <ف> for /v/ or appending a small *v* below or above: <ڤ> for /r/, and <ڭ> for /ɣ/. Similar to other Perso-Arabic writing systems, the vowel set borrows from the consonant set in that *yeb* and *waw* double as vowels and consonants. *Alef* is used as a vowel. The long /u:/ is indicated by doubling the *waw*. *Waw* and *yeb* with a small *v* above indicate /o/ and /e/, respectively.

Although the *kasra* is not part of the modern Sorani orthography, it is rarely used in some dictionaries for disambiguating certain pronuncia-

22. <https://www.ethnologue.com/language/ckb>

23. In 2006, Duhok Governorate began using Kurmanji as their official language as a way of resisting Sorani.

24. <https://sinaahmadi.github.io/resources/pewan.html>

tions,²⁵ where it is used to mark a short /I/ that is otherwise unrepresented in modern Perso-Arabic orthography.²⁶

3.6. Uyghur

Uyghur, also written as Uighur, is a Turkic language spoken in the region in and around what is known as the Xinjiang Uyghur Autonomous Region in Northwest China.²⁷ Due to a politically and culturally motivated diaspora, Uyghur is spoken in Turkic countries such as Kazakhstan, Kyrgyzstan, Uzbekistan and Turkey, but also by the smaller Uyghur migrant communities elsewhere (Dillon, 2009). Ethnologue estimates the number of Uyghur speakers at around 10 million in China and total speakers in all countries at around 10.4 million. The modern Uyghur writing system should not be confused with Old Uyghur which was written using Sogdian script (Wilkins, 2016). Historically the writing system dates back to the 10th century when the Perso-Arabic script was introduced along with the spread of Islam and which evolved after considerable changes over the centuries into what is recognized as the modern Uyghur Perso-Arabic orthography (Brose, 2017). The writing system for the language underwent extensive changes, including being changed to the Cyrillic and Latin scripts and even the Pinyin romanization system for political reasons. It was not until 1982 that the Arabic Uyghur alphabet was reinstated (Dwyer, 2005). As of today the language has four writing systems: Uyghur Arabic used in the Xinjiang province of China, Uyghur Cyrillic in Kazakhstan, Uyghur Latin in Turkey and Uyghur Pinyin, which is not used much (Hamut and Joniak-Lüthi, 2015).

Unlike most other writing systems using Perso-Arabic, but like Sorani, Uyghur writing system is an alphabet, i.e., the vowels are explicitly marked. Uyghur is written in Naskh style, a sample of which is shown in Table 7.²⁸ As is the case with all languages which have adopted the Perso-Arabic script, in order to represent the phonemic features of the language, Uyghur writing system has evolved an original repertoire of letters. Apart from the letters borrowed from the original Arabic script, four letters are derived from Persian writing system: <پ>, <چ>, <ژ>,

25. Private correspondence from Aso Mahmudi (2022).

26. According to Ahmadi (2019, §2.2, p.3), the corresponding letter of Latin-based orthography of Kurmanji dialect is ⟨i⟩.

27. <https://www.ethnologue.com/language/uig>

28. The Latin text obtained from UygurAvazi newspaper (<https://uyguravazi.kazgazeta.kz/>) was converted to Uyghur through a script converter from <http://www.elipbe.com>.

TABLE 7. Uyghur sample in Naskh transliterated from Latin text

Naskh
پره زېدهنت ماراسم قاتناشقۇچىلىرىنى ۋاتان قىمايچىسى كۈنى ۋا غالىبىيات كۈنى بىلەن تاپرىكلاپ قارىيى كەزىماتچىلارنىڭ مىنلاشكا ئالاقىدا قاسسا قوشۇۋاتقانلىغىنى ئاتاپ كورساتتىن
Latin
prezident marasim qatnashquchilirini vatan qimayichisi kuni va ghalibiyat kuni bilan tabriklap qarbiy khizmatchilarninh taminlashka alaqlida qassa qoshuvatqanlighini atap korsattin

TABLE 8. A line from South Azerbaijani Wikipedia: “Mirza Shafi Vazeh—Azerbaijani poet, thinker, enlightener and teacher”

Nastaliq
میرزا شفیع واضح آذربایجان شاعیری و موثقیری، معارفچی و پداقوق.
Naskh
میرزا شفیع واضح آذربایجان شاعیری و موثقیری، معارفچی و پداقوق.

<گ>; the <ك>, which represents a velar nasal, common to Turkic languages, is derived from the Arabic *kaf* <ك> with three dots positioned above the letter. The *two-eyed he* <ه> is also used, similar to Kazakh, Urdu, Sindhi and Shahmukhi among other languages.

Extensive use of the *waw* is made, which is modified in productive ways to represent the vowels: <ؤ> with a superscript *alef* to represent the sound /yu/, <ۇ> with a *small v* on top to represent a front rounded vowel /ø/, <ۇ> with a *damma* on top for a long /u:/ and <ۇ> with *three dots* on top represents the semivowel /w/. Uyghur uses the Arabic *yeb* <ي> for the semivowel /j/, the *alef maqsura* <ى> for the /i/ and <ې> *yeb* with two dots below to represent /e/. Additional combining marks are used to mark consonant clusters (Arabic *sukun*, U+0652) or gemination (Arabic *shadda*, U+0651).

3.7. Southern Azerbaijani

Azerbaijani, also known as Azeri, Azari, Azeri Turkish and Azerbaijani Turkish, belongs to the Turkic language family, more specifically to the Western Oghuz branch (Mokari and Werner, 2017). It is spoken by over 23 million people, mainly in Azerbaijan, Iran, Georgia, Rus-

sia and Turkey, and also in Iraq, Syria and Turkmenistan.²⁹ Two varieties of the language are recognized: Northern Azerbaijani and Southern Azerbaijani. Northern Azerbaijani is spoken in the Republic of Azerbaijan, where it is the official language. Southern Azerbaijani spoken by around 14.6 million people is confined to the northwest of Iran and is often called *Turki* (تورکی).³⁰ Due to migrations and trade it is also used in parts of Iraq and Turkey, and in Afghanistan and Syria. Whereas Northern Azerbaijani uses either the Cyrillic script (in Dagestan) or Latin (the official script in Azerbaijan), Southern Azerbaijani uses the Perso-Arabic script. The Naskh style is favoured in day to day use but Nastaliq is sometimes used, mainly for book titles and also for handwriting, as demonstrated in Table 8.

Historically Old Azeri (*Ādari*) was the Indo-Iranian language spoken in Persian Azerbaijan before the arrival of the Turkic-speaking populations to the region (Yarshater, 2011). The language was gradually replaced with Turkish as the migration of Turkic speakers increased and by the 18th century Turkish was recognized as the language of Azerbaijan, although the name *Ādari* was retained as Azeri and traces of Old Azeri can still be found in Turkish today (Bosworth, 2011). The arrival of the muslim Turkish speakers in South Azerbaijan was accompanied by the Perso-Arabic *abjad* which became the official script of Azerbaijan until the 1920's, when, for political reasons, competing Cyrillic and Latin scripts entered the scene (Hatcher, 2008). The Azerbaijani Perso-Arabic writing system saw considerable mutation over the centuries: 28 letters (all from Arabic) initially, increased to 32 letters with additions from Persian and, finally, 33 letters due to an addition from the Ottoman Turkish. None of these solutions were found suitable for Azerbaijani and reforms were proposed during the 19th and 20th centuries which finally created the character set of Southern Azerbaijani as it is known today.³¹

The modern inventory consists of 42 letters. The majority of letters are borrowed from the Arabic and Persian orthographies. Nine letters (<ذ>, <ژ>, <ص>, <ض>, <ط>, <ظ>, <ع>, <ح>, and <ث>) are exclusively used for spelling Persian and Arabic loanwords and names. An extra letter *kebeh with three dots above* <ک̣> is used to indicate the voiced velar nasal /ŋ/, similar to Uyghur (Daniels, 2014, p. 31). Like all Turkic languages, Azerbaijani has a rich vowel system (Johanson and Csató, 2021). Three core shapes <ا>, <و>, and <ی>, modified with various diacritics, form the letters of the vowel set. Letter <ئ> represents the sound /e/, <ئ> the unrounded back vowel /u/ and <ي> represents /i:/. The *rasm* for *waw* is adapted in four ways. Apart from the intrinsic value of

29. <https://www.ethnologue.com/language/aze>

30. <https://www.ethnologue.com/language/azb>

31. For example, in modern Azerbaijani, letter *kebeh* <ک̣> has replaced the older Arabic *kaf* <ک>.

TABLE 9. Sample of Jawi in Naskh (top) and Rumi (bottom)

Naskh
ينله كراغن سواتو ماده مغارغن شعير نرلاو اينده، ممبتولي جالن تمقت برقيبنده، د سانله ايتيكات دقريتولي سوده
Rumi
inilah gerangan suatu madah mengarangkan syair terlalu indah, membetuli jalan tempat berpindah, di sanalah i'tikat diperbetuli sudah

waw, <ؤ> is used for /u:/, <ؤ> for the front open rounded vowel /y/, <ؤ> for the front rounded vowel /ø/ and the digraph *waw with sukun* <ؤ> represents /o/. The letter <ه> is used only in the final form to mark the diphthong /ae/. In addition, Azerbaijani orthography admits three combining marks, *fatha*, *damma* and *kasra*, to mark short vowels, and also additional diacritics to mark consonant clusters (*jazm/sukun*, U+0652) or gemination (*shadda*, U+0651).

3.8. Malay (Jawi)

Jawi is a Perso-Arabic writing system used for recording the Malay language from the Austronesian family and several other languages of Southeast Asia (Kratz, 2002). With the advent of Islam in Southeast Asia around the 14th century, the Pallava script, Nagari, and old Sumatran scripts which were used in writing Malay, were replaced by the Perso-Arabic script and by the 15th century Jawi had spread to Brunei, Indonesia and even Thailand due to trading (Coluzzi, 2020). Its dominance remained till the 20th century when Jawi was replaced by the Latin script (*Rumi*) and was confined to religious and cultural rituals. Today apart from Malaysia, Jawi has the status of an official writing system in Brunei and also in Indonesia, where Jawi has been assigned a regional status (Abdullah et al., 2020). Unlike Urdu or Shahmukhi, and, to a lesser extent Persian, Jawi favours the Naskh style, demonstrated by the sample quatrain from “Syair Perahu,” a Sumatran Sufi poem (Braginsky, 1975), in Table 9.

In addition to the 28 basic characters from Arabic,³² Jawi added extra characters to suit its requirements and introduced the following: <چ> *ca*,

32. Some scholars, like R. O. Windstedt, believe that Jawi borrowed the characters from the Persian, rather than Arabic, orthography (Windstedt, 1961).

<ڠ> *nga*, <ڤ> *pa/fa*,³³ <ڭ> *ga*, and <ن> *nya*. The letter <ڤ> *v* was added for representing foreign loanwords.³⁴ This brings the size of modern Jawi inventory to 37 letters (DBP, 2006). In addition, three more characters are possible due to the adjunction of the *high hamza* <ء> above *alef* <ا>, below *alef* <اِ>, and above *yeh* <ي> (MS, 2012).

As in Arabic, vowel diacritics are not normally used. However, if needed Jawi has three diacritics used to indicate the short vowels: *fatha*, *damma*, and *kasra*. A major feature of the language is the use of full reduplication of the base word (Prentice, 1990). This is represented in Jawi with the Arabic numeral <٢> (“2”) as in انجڠ-انجڠ *anjeng-anjeng* (“dogs”) as a shorthand for the equivalent longer spelling انجڠ-انجڠ for the plural form of a noun انجڠ /andʒɛŋ/ (“dog”).

4. Finite-state Transformations of Perso-Arabic Script

Below we provide a brief overview of the design for Perso-Arabic script normalization framework provided by the open-source Nisaba software package.³⁵ The design was partially inspired by prior formal approaches to computational modeling of Brahmic alphasyllabaries (Datta, 1984; Sproat, 2003) and, in particular, our prior work at Brahmic script normalization (Gutkin, Johny, Doctor, Wolf-Sonkin, et al., 2022; Johny, Wolf-Sonkin, Gutkin, and Roark, 2021). These approaches exploit the inherent structure which manifests itself in all the Brahmic *abugidas* in the notion of “orthographic syllable” or *akṣara* (Bright, 1999; Fedorova, 2012). In contrast to various Brahmic scripts, the Perso-Arabic *abjad* does not offer the same rigid orthographic structure. Nevertheless, a similar in nature formal approach to script normalization, designed to address the kind of Perso-Arabic script representation ambiguities outlined in §2, can be pursued. We previously showed that scripts, such as Thaana, that borrow their features from both script families are amenable to such formal analysis (Gutkin, Johny, Doctor, Wolf-Sonkin, et al., 2022).

Our script processing pipeline consists of multiple components implemented as finite-state grammars using Pynini (Gorman, 2016; Gorman and Sproat, 2021), which is a Python framework for compiling grammars expressed as strings, regular expressions, and context-dependent rewrite rules into (weighted) finite-state transducers (FSTs).

33. The letter *fa* <ڤ> was used to represent *pa* because the sound /f/ does not exist in Malay and was pronounced as /p/.

34. The letter *va* <ڤ> is mostly used to spell English loanwords, e.g., اونيفرسيتي (“universiti”).

35. For more detailed treatment of this software please see Gutkin, Johny, Doctor, Roark, et al. (2022).

TABLE 10. Summary of script transformation operations

Operation Type	FST	Language-dependent	Includes
NFC	\mathcal{N}	no	—
Common Visual	\mathcal{V}_c	no	\mathcal{N}
Visual	\mathcal{V}	yes	\mathcal{V}_c
Reading	\mathcal{R}	yes	\mathcal{V}

The resulting FSTs can then be efficiently combined together in a single pipeline in a variety of downstream applications (Mohri, 1996; 2009). These component FSTs are shown in Table 10 and described below.

Unicode Normalization

There exist language-agnostic procedures—part of the Unicode standard—that normalize text with Perso-Arabic string encodings to visually equivalent canonical normal forms. Normalization Form C (NFC) is a well-known and widely-used standard of this sort, and its application results in an equivalence class of visually identical strings that are all mapped to a single conventionalized representative of the class (Whistler, 2021). In the Nisaba library, the NFC standard is operationalized by compiling the transformations into an FST, which we denote as \mathcal{N} in Table 10. The transformations include compositions and re-orderings, along with combinations of multiple such transformations.

Composition transformations can be illustrated with the following concrete example. The *alef with madda above* letter <|> has two visually identical possible encodings: with two characters by adjoining *maddah above* to *alef* ($\{ \text{U+0627}, \text{U+0653} \}$), or as the single character that already includes the maddah (U+0622). The FST \mathcal{N} transforms the two character encoding into the single character encoding, which does not change the appearance of the letter. Re-ordering transformations address multiple encodings that can arise with Arabic combining marks. As a concrete example, *shadda* (U+0651) followed by *kasra* (U+0650) yields the same rendering as *kasra* (U+0650) followed by *shadda* (U+0651). The NFC canonical form is the latter, hence the \mathcal{N} FST transforms the former encoding to the latter. The string $\{ \text{alef (U+0627)}, \text{superscript alef (U+0670)}, \text{maddah above (U+0653)} \}$ is an example that transforms via \mathcal{N} with both composition and re-ordering to the visually identical form $\{ \text{alef with madda above (U+0622)}, \text{superscript alef (U+0670)} \}$.

As noted above, \mathcal{N} is language-agnostic, meaning that its transformations (taken from the NFC standard) do not violate any language’s writing system rules.

Visual Normalization

We use the term *visual* normalization—initially introduced in the context of Brahmic script normalization (Johny, Wolf-Sonkin, Gutkin, and

TABLE 11. Example Urdu components included in language-specific FST \mathcal{V}_l

Kind of rewrite	FST	Letter	Variant (source)	Canonical (target)
position-independent	\mathcal{V}_l^*	$\langle j \rangle$	<i>reb + small high tab</i>	<i>rreb</i>
non-final	\mathcal{V}_l^n	$\langle \mathcal{S} \rangle$	<i>kaf</i>	<i>kebeb</i>
word-final	\mathcal{V}_l^f	$\langle \mathcal{S} \rangle$	<i>alef maksura</i>	<i>farsi yeb</i>
isolated-letter	\mathcal{V}_l^i	$\langle o \rangle$	<i>beb</i>	<i>beb goal</i>

Roark, 2021)—to denote transformations that are not part of NFC but that also result in canonical forms that are visually identical to the input. This is implemented via two FSTs, one for language-agnostic transformations and one language-specific, which are combined (via FST composition) with NFC into a single language-dependent FST: $\mathcal{V} = \mathcal{N} \circ \mathcal{V}_c \circ \mathcal{V}_l$, where \circ denotes the composition operation (Mohri, 2009).³⁶

The language-agnostic FST, \mathcal{V}_c , consists of the small set of normalizations not included in NFC that apply to all supported languages. As a concrete example of this class of transformations, the two-character encodings of *warw* (U+0648) followed by either *damma* (U+064F) or *small damma* (U+0619) are mapped to *u* (U+06C7). Perso-Arabic “presentation forms” from Unicode Block A, which include ligatures and contextual forms for letter variants required by the writing systems for Persian, Urdu, Sindhi and Central Asian languages,³⁷ are also normalized to visually identical canonical forms by \mathcal{V}_c , as specified by Unicode NFKC normalization (Whistler, 2021). For example, letter *beeb isolated form* $\langle \mathcal{P} \rangle$ (U+FB52) is normalized to *beeb* (U+067B), which is visually identical. The character *ligature lam with alef isolated form* $\langle \mathcal{L} \rangle$ (U+FEFB) is transformed to two characters: *lam* $\langle j \rangle$ (U+0644) followed by *alef* $\langle i \rangle$ (U+0627).

Language-specific visually-invariant transformations, included in the FST denoted as \mathcal{V}_l , include four special cases related to positions in the word: word-final, non-final (i.e., word-initial and word-medial), isolated-letter and position-independent transformations. Each of these are compiled into their own FST, as shown in Table 11, then composed into a single $\mathcal{V}_l = \mathcal{V}_l^i \circ \mathcal{V}_l^f \circ \mathcal{V}_l^n \circ \mathcal{V}_l^*$. Table 11 additionally presents some example transformations of each type, taken from the set of transformations required for Urdu.

Reading Normalization

Gutkin, Johny, Doctor, Wolf-Sonkin, et al. (2022) noted the need for some additional normalization beyond those preserving visual identity for the Brahmic scripts, which they termed *reading* normalization.

36. Johny, Wolf-Sonkin, Gutkin, and Roark (2021) provides details regarding composition and other operations used by FSTs in these normalizers.

37. <https://www.unicode.org/charts/PDF/UFB50.pdf>

We also include this class of normalizations for Perso-Arabic, which we compile into the FST denoted \mathcal{R} in Table 10. Full reading normalization is the finite-state composition of visual normalization with language-specific reading normalization: $\mathcal{R} = \mathcal{V} \circ \mathcal{R}_l$. For example, Persian, Shahmukhi, Kashmiri, Urdu and Sorani Kurdish all map from *yeb* $\langle \text{ﻲ} \rangle$ (U+064A) to *farsi yeb* $\langle \text{ﻲ} \rangle$ (U+06CC), while Uyghur, Sindhi and Malay employ the inverse of this transformation, as dictated by their respective orthographies.

5. Experiments

While outlining the potential issues that may arise with text written in the Perso-Arabic script is important, it is also useful to assess how common the issues may be in real-world text. To that end, we devised some experiments that derive natural language models from collected text and validate their quality both with and without normalization. If the phenomena being normalized are rare, then the difference between the conditions will be small; and if the normalizations do not result in better text representations, then the normalized conditions may exhibit a lower quality in the validation. In this section, we present the details of our assessment, first for statistical language modeling, which provides an intrinsic validation of model quality, followed by machine translation, which provides an extrinsic validation. As was mentioned in the introduction, the code for the experiments and the corresponding results for both validation types have been released (see §1 on page 317).

5.1. Language Modeling Experiments

For language modeling experiments we use Wikipedia data for eight languages: five Indo-European—Kashmiri, Kurdish (Sorani), Punjabi, Sindhi, and Urdu; Malay from an Austronesian group; and Uyghur and Azerbaijani from the Turkic group. A brief overview of experimental methodology is given in §5.1.1. The dataset preprocessing details are provided in §5.1.2. The details and results of statistical language modeling experiments can be found in §5.1.3.

5.1.1. Methodology

Language models are trained to predict the next token in a sequence given the previous tokens. Tokens can be variously defined as characters or words, or even as morepheme-like sub-word multi-character tokens. The intrinsic quality of the language model can be measured via the probability the language model assigns to attested exemplars, i.e.,

real text. The higher the aggregated probability of the attested text, the better the language model. See, e.g., Rosenfeld (2000) for more details on this long-standing validation paradigm. For this work, a key consideration is *comparability*—we need to ensure that models have access to the same training data and that the validation data is identical.

To ascertain whether script normalization has any significant impact on language model quality we follow a simple methodology. We adopt a k -fold cross-validation design where, for each language, we randomly shuffle the dataset and split it into the 80% training and 20% test folds, repeating the process k times, where $k = 100$. At each iteration we train the models (e.g., language models as in §5.1.3) and evaluate them by computing the corresponding metric (e.g., cross-entropy).

Following the above procedure, statistics are assembled using k observations for the baseline configurations that correspond to the original text and the actual testing configurations corresponding to the normalized text. Crucially, we start by generating the normalized data, recording all the sentences which contain the actual diffs in set D . During the generation of the training and test data for the baseline and testing configurations, we make sure that the sentences in D are confined to the training set. In other words, we make sure that all the actual rewrites are confined to the training data for all the k folds, which ensures that the test folds are always identical for normalized and unnormalized conditions. Given the baseline and the test metric distributions, we employ significance testing to validate the null hypothesis that the two distributions are identical; in other words, that the normalization has no significant impact on the model performance.

Three types of statistical hypothesis tests are used here. Assuming that the two groups are normally distributed an obvious choice is the two-sample (independent) t -test for comparing the means of the two populations (Zabell, 2008). Making an additional assumption that the population variances are not equal, we employ Welch’s formulation of t -test (Welch, 1947) with Satterthwaite’s degrees of freedom (Satterthwaite, 1946), referred to below as Welch-Satterthwaite (WS) test. The test provides the t statistic, the p -value and the estimated confidence interval (CI) $[L, H]$ for the 95% confidence level at the significance level of $\alpha = 0.05$.

In addition, two non-parametric approaches are used here. A Mann-Whitney (MW) test (Mann and Whitney, 1947) and a more recent Brunner-Munzel (BM) test (Brunner and Munzel, 2000). Both tests provide the t statistic and the p -value. The rationale for using multiple

TABLE 12. Details of preprocessed Wikipedia datasets

Language	Code	β	N_l	N_w	N_l'	$R_l(\%)$	N_w'	$R_w(\%)$
Kashmiri	ks	0.6	3266	9721	530	16.22	442	4.55
Kurdish (Sorani)	ckb	0.8	839 750	794 475	42 437	5.05	67 569	8.5
Malay	ms	0.0	102 311	200 052	72 645	71.0	25 854	12.92
Punjabi (Shahmukhi)	pnb	0.8	1 075 820	886 399	145 391	13.51	16 443	1.86
Sindhi	sd	0.8	201 345	240 591	138 839	68.96	64 006	26.6
South Azerbaijani	azb	0.8	1 638 622	735 986	60 094	3.67	23 834	3.32
Uighur	ug	0.1	110 344	376 307	3600	3.26	19 461	5.17
Urdu	ur	0.9	3 595 095	799 610	6600	0.18	3632	0.45

hypothesis tests is to see whether they all agree with other providing additional weight to the null (or alternative) hypothesis.³⁸

5.1.2. Corpus Preprocessing

The process of preparing the Wikipedia data is kept simple. The datasets for each language are downloaded in the MediaWiki XML format. The particular version of the dump is restricted to the pages with their current versions including the metadata.³⁹ The key difficulty lies in extracting the actual plain text in native language from the structured XML data while weeding out the metadata. We use the `mwxml` Python package developed by the Wikipedia foundation to iterate over the articles in MediaWiki XML dump.⁴⁰

For each article, we use the MediaWiki Parser from Hell package to parse the current revision of article's text.⁴¹ Once the parse is complete, we strip the contents of all the "unprintable" content, such as templates, using the API provided by the `mwparserfromhell` package, and split the text by newlines. A simple script detection and filtering algorithm is used to decide whether to keep the sentence or drop it from the resulting

38. All the algorithms are provided by the open-source <https://docs.scipy.org/doc/scipy/reference/stats.html> `scipy.stats` and <https://www.statsmodels.org/stable/index.html> `statsmodels` packages.

39. For example, a reasonably recent dump for Punjabi (Shahmukhi) is available at <https://dumps.wikimedia.org/pnbwiki/20211120/pnbwiki-20211120-pages-meta-current.xml.bz2>.

40. https://www.mediawiki.org/wiki/Mediawiki_utilities/mwxml

41. <https://github.com/earwig/mwparserfromhell>

data.⁴² Any given l -character long sentence is dropped from the data if it contains less than $\beta \cdot l$ characters in native (Perso-Arabic) script, where $\beta \in [0, 1)$. The filtering factor β is language-specific and is determined by informally examining the data.⁴³ This filtering process is crude in that it excludes any control for sentence or token length.

The preprocessing details are shown in Table 12. Each language is shown along with its Wikipedia code, the script filtering factor β described above, the number of resulting lines N_l and the number of unique tokens N_w . For the corresponding normalized text, N_l^r denotes the number of lines that contain diffs and N_w^r is the number of unique tokens that differ from the unnormalized version. The ratios between modified lines and token types are denoted R_l and R_w , respectively. The tokenization process is relatively crude and involves splitting on the whitespace completely disregarding other types of punctuation, such as Perso-Arabic punctuation symbols. We refer to the output of tokenization as tokens, rather than words, because the data is quite noisy even after filtering.

When normalizing text we apply the Nisaba *reading* normalization grammar (Gutkin, Johny, Doctor, Wolf-Sonkin, et al., 2022), which subsumes all the grammars providing visual invariant transformations, i.e., NFC and *visual* normalization (Johny, Wolf-Sonkin, Gutkin, and Roark, 2021), as well as transformations that change the visual appearance of the Perso-Arabic tokens. According to Table 12, the normalization effects vary across languages. For Urdu, which is the largest dataset, the percentage of modified lines and token types is below one percent. This reflects the relatively low number of transformations currently enabled in the Nisaba Urdu grammars compared to the other six languages. The highest proportion of modified lines and tokens happens in Sindhi and Malay, while for Kashmiri (the smallest datasets) and Punjabi (Shahmukhi, the second largest) the number of modifications is relatively low. A description of how these modifications affect model quality follows next.

5.1.3. Statistical Language Models

For building n -gram language models we use the KenLM toolkit (Heafield, 2011)⁴⁴ which is fast and easy to use relative to alternatives. We used modified Kneser-Ney modeling options, as recommended (Heafield, Pouzyrevsky, Clark, and Koehn, 2013b). In what follows, the terminology introduced in §5.1.1 is used. The experiments with

42. We previously implemented a similar script detection algorithm for the Wikipron project <https://github.com/CUNY-CL/wikipron>.

43. For Malay we use $\beta = 0$, i.e., no filtering.

44. <https://github.com/kpu/kenlm>

TABLE 13. Character-level statistics for k -fold configurations and m n -gram orders ($k = 100$, $m = 8$)

Language	Train		Test	
	μ	σ	μ	σ
azb	148 133 643.5	104 532.9	32 841 396.5	104 532.9
ckb	141 088 021.3	118 118.6	32 419 200.7	118 118.6
ks	252 357.7	3438.0	63 370.3	3438.0
ms	16 981 259.9	9569.6	3 394 367.1	9569.6
pnb	216 833 075.2	248 147.7	54 209 912.8	248 147.7
sd	36 855 104.9	95 212.8	9 212 888.1	95 212.8
ug	33 813 243.4	109 518.9	8 096 172.6	109 518.9
ur	368 794 342.1	240 334.8	92 205 182.9	240 334.8

character and word n -gram models are described in §5.1.3 and §5.1.3, respectively.

For a single fold, the criterion for splitting into training and test sets is to use the number of lines in the corpus. As a result, the number of training and test tokens (whether these are individual characters or words) differ across the folds.

Character Models

For each language and each of the $k = 100$ train/test folds we build n -gram character language models for orders $n \in [3, \dots, 10]$. The character-level statistics computed for each language over all the folds and all the orders (amounting to 800 observations per language) are shown in Table 13, where the means and standard deviations are shown for the training and test datasets. Since the corpora are split by the number of lines, the resulting variances for character datasets are quite high.

The resulting cross-entropies (in bits per character) for the models built in this way from the unnormalized text are shown in Figure 2, where each point for each n -gram order in the curve is shown along with its corresponding error band computed over k models. The plot for Kashmiri, the smallest dataset among the four languages, stands out in that the error band is clearly visible, especially for the higher orders for which the model overfits the training data. The plots for the rest of the languages show very low variance at each point in the plot for all the n -gram orders.

For each of the languages and each of the n -gram orders, we perform statistical hypothesis testing for the differences in mean cross-entropies between the character language models trained on the original (baseline, denoted B) and the normalized (test, denoted T) text for all the k folds ($k = 100$). As mentioned in §5.1.1, for each fold, the sentences that contain (for T) or act as source of (for B) normalization diffs are kept in the training portion of the data. Full results for each of the languages




FIGURE 2. Average values of entropy (bits per character) over 100 runs vs. character n -gram orders

are presented in Appendix B, and key summarizing values are shown for all in Table 14.

The mean cross-entropy difference Δ_μ is computed over all the k folds as

$$\Delta_\mu = \frac{1}{k} \sum_{i=1}^k (H_i(T_i) - H_i(B_i)), \quad (1)$$

where the negative value of Δ_μ indicates the decrease of character entropy H of model i compared to the baseline and hence constitutes an improvement.

As can be seen from the table, the Δ_μ values are negative across the board, apart from the lowest n -gram orders ($n = 3$ and $n = 4$) for Sindhi. To determine whether these changes in cross-entropy are statistically significant, three types of tests (WS, MW and BM) were performed (see §5.1.1). All of the tests assess the null hypothesis that baseline and test configurations represent the same distribution. While we do not explicitly compute the correlation between the p -values for all the three tests, these tend to correlate with each other upon informal inspection. All significance test values for all languages are presented in Appendix B. Since the trends are largely the same, for ease of inspection we just show the WS p -value in Table 14, where the statistically significant degradation for Sindhi configuration corresponding to $n = 4$ is marked in red, and discuss the few disagreements in the Appendix.

TABLE 14. Significance tests for character n -gram language models. Δ_μ is the mean absolute change in cross-entropy after normalization; % is the percentage change; and p is the WS p -value.

Language	Measure	n -gram order							
		3	4	5	6	7	8	9	10
South Azerbaijani	Δ_μ	-0.012	-0.004	-0.002	-0.001	-0.002	-0.002	-0.002	-0.002
	%	0.418	0.186	0.131	0.093	0.14	0.143	0.13	0.171
	p	0.0	0.0	0.002	0.048	0.012	0.013	0.02	0.004
Kurdish (Sorani)	Δ_μ	-0.01	-0.003	-0.004	-0.003	-0.004	-0.006	-0.005	-0.005
	%	0.36	0.16	0.24	0.2	0.28	0.44	0.43	0.39
	p	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Kashmiri	Δ_μ	-0.003	-0.021	-0.006	-0.006	-0.005	-0.016	-0.022	-0.028
	%	0.08	0.71	0.21	0.22	0.2	0.63	0.85	1.07
	p	0.647	0.007	0.544	0.515	0.564	0.075	0.014	0.003
Malay (Jawi)	Δ_μ	-0.065	-0.062	-0.06	-0.064	-0.067	-0.07	-0.07	-0.07
	%	1.818	2.036	2.232	2.526	2.736	2.885	2.931	2.922
	p	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Punjabi (Shahmukhi)	Δ_μ	-0.011	-0.013	-0.01	-0.008	-0.007	-0.006	-0.007	-0.007
	%	0.32	0.44	0.39	0.34	0.33	0.3	0.36	0.34
	p	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Sindhi	Δ_μ	0.001	0.011	-0.018	-0.028	-0.016	-0.024	-0.015	-0.014
	%	-0.03	-0.28	0.47	0.79	0.46	0.71	0.45	0.42
	p	0.479	0.0	0.001	0.0	0.001	0.0	0.007	0.015
Uyghur	Δ_μ	-0.002	-0.001	-0.004	-0.004	-0.003	-0.004	-0.004	-0.005
	%	0.074	0.051	0.203	0.219	0.164	0.225	0.24	0.295
	p	0.0	0.025	0.0	0.0	0.0	0.0	0.0	0.0
Urdu	Δ_μ	-0.003	-0.005	-0.004	-0.004	-0.004	-0.002	-0.002	-0.005
	%	0.11	0.22	0.2	0.26	0.24	0.14	0.13	0.37
	p	0.0	0.0	0.0	0.0	0.0	0.02	0.217	0.001

TABLE 15. Word-level statistics for k -fold configurations and m n -gram orders ($k = 100$, $m = 4$)

Language	Train		Test	
	μ	σ	μ	σ
azb	9 704 485.4	7010.9	2 110 255.6	7010.9
ckb	10 462 576.0	9503.5	2 384 669.0	9503.5
ks	23 189.4	250.5	4164.6	250.5
ms	1 463 801.4	814.4	290 574.6	814.4
pnb	24 690 883.4	13 213.2	3 111 613.6	13 213.2
sd	4 680 624.1	662.0	74 124.0	662.0
ug	2 160 932.7	7224.1	515 917.3	7224.1
ur	37 234 659.5	23 873.1	9 235 049.5	23 873.1




FIGURE 3. Average values of entropy (bits per word) over 100 runs vs. word n -gram orders

Word Models

We repeated all the experiments described in §5.1.3 above for the n -gram models trained on words for the orders 2, 3, 4 and 5. The details of the training and test splits are shown in Table 15. As mentioned above, the Kashmiri dataset is very small and this sparsity is only increased when considering word-sized tokens instead of characters (compare this with Table 13). This is reinforced by computing the word cross-entropies for Kashmiri models, shown in Figure 3, where the error band shows significantly higher variance (compared to the character models in Figure 2) across each k splits for all the orders compared to other languages. Sindhi and Malay, which are the second and third smallest datasets, show a reasonably high variance as well, although it is significantly smaller than for Kashmiri. The plot for Kurdish (Sorani) indicates that the quality of the word models tends to degrade for this corpus beyond trigrams, possibly due to a relatively small size of the dataset.

Statistical significance tests were also performed for the word n -gram models constructed for $n \in [2, 3, 4, 5]$ from the $k = 100$ folds over original and normalized text. All values for all languages are presented in Appendix B, and key measures over all languages are shown in Table 16. Again, we just show the WS p -value in this summary table, but the values for all tests are presented and discussed in the Appendix. Kashmiri results are not significant for any of the orders, likely due to the very small size of the dataset. All other languages show statistically significant reductions in cross-entropy for all n -gram orders. Reductions are relatively small for Kurdish, Punjabi, Azerbaijani, Uyghur and Urdu, but

TABLE 16. Significance tests for word n -gram language models. Δ_μ is the mean absolute change in cross-entropy after normalization; % is the percentage change; and p is the WS p -value

Language	Measure	n -gram order			
		2	3	4	5
South Azerbaijani	Δ_μ	-0.031	-0.027	-0.031	-0.028
	%	0.374	0.332	0.397	0.358
	p	0.0	0.0	0.0	0.0
Kurdish (Sorani)	Δ_μ	-0.031	-0.034	-0.034	-0.035
	%	0.41	0.49	0.49	0.5
	p	0.0	0.0	0.0	0.0
Kashmiri	Δ_μ	0.034	-0.039	-0.022	0.017
	%	-0.33	0.39	0.22	-0.17
	p	0.112	0.127	0.34	0.468
Malay (Jawi)	Δ_μ	-0.358	-0.394	-0.403	-0.411
	%	2.935	3.319	3.411	3.479
	p	0.0	0.0	0.0	0.0
Punjabi (Shahmukhi)	Δ_μ	-0.015	-0.017	-0.02	-0.02
	%	0.15	0.19	0.23	0.23
	p	0.0	0.0	0.0	0.0
Sindhi	Δ_μ	-0.159	-0.169	-0.177	-0.184
	%	1.06	1.21	1.26	1.32
	p	0.0	0.0	0.0	0.0
Uyghur	Δ_μ	-0.024	-0.038	-0.034	-0.042
	%	0.2	0.332	0.294	0.367
	p	0.0	0.0	0.0	0.0
Urdu	Δ_μ	-0.006	-0.004	-0.006	-0.007
	%	0.08	0.06	0.1	0.1
	p	0.0	0.001	0.0	0.0

more substantial for Sindhi and particularly Malay, which achieves up to 3.5% reduction in cross-entropy.

In sum, we have shown consistently small yet significant improvements in intrinsic language model quality through the use of these normalization methods.

5.2. Neural Machine Translation Experiments

This section describes the application of Perso-Arabic script normalization to machine translation (MT), which is arguably one of the oldest and most popular downstream NLP tasks (Hutchins, 1986). We selected a subset of languages described in §3 and designed a simple translation experiment, where for each language we build a model that translates that source language into English in two configurations: the model trained on the original source text and the model trained on the normal-

TABLE 17. Parallel corpora details for the four languages. The dataset sizes for each language correspond to the total number of parallel sentence pairs in all datasets for a particular language.

Languages	Train	Development	Test
Kurdish	XLENT (ckb), TICO-19 (ckb), Wikimedia (ckb), OPUS-100 (kur), TANZIL (kur), TATOEBEA (kur) 256,909 sent. pairs	OPUS-100 (kur) 2,000 sent. p.	OPUS-100 (kur) 2,000 sent. p.
Sindhi	CCMATRIX, XLENT, TANZIL, QED, Wikimedia 1,960,022 sent. pairs	Ubuntu 6,204 sent. p.	CCMATRIX (held-out) 1,000 sent. p.
Urdu	OPUS-100, JOSHUA, ANUVAAD 798,574 sent. pairs	OPUS-100 2,736 sent. p.	OPUS-100 2,000 sent. p.
Uyghur	XLENT, TANZIL, TATOEBEA, TED, OPUS-100 176,179 sent. pairs	OPUS-100 2,000 sent. p.	OPUS-100 2,000 sent. p.

ized source text. We hypothesize that if the normalization is “useful,” it will result in a better model of the source language (by removing the extrinsic orthographic artifacts of Perso-Arabic resulting in systematic ambiguities) and, consequently, a better translation quality into English as measured by the objective evaluation metrics.

In what follows we introduce the parallel language corpora used for training and evaluating the models in §5.2.1, provide brief summary of the monolingual and multilingual model architectures used in §5.2.2, and discuss our results in §5.2.3. It is important to note that our aim here is not to produce a competitive MT system using current state-of-the-art (such as Wenzek et al., 2021; Xue, Barua, et al., 2022; Xue, Constant, et al., 2021), but rather to measure the effects of script normalization using reasonably advanced yet simple-to-train neural models.

5.2.1. *Parallel Corpora*

In our experiments we construct individual models for translating from four languages into English: Kurdish, Sindhi, Urdu and Uyghur. These parallel corpora were collected using the MTDATA tool that automates the collection and preparation of machine translation datasets (Gowda, Zhang, C. Mattmann, and May, 2021).⁴⁵ Each language may have multiple datasets available from several sources, such as the OPUS collection that provides various machine translation corpora for many lan-

45. <https://github.com/thammegowda/mtdata>

guages (Tiedemann, 2012).⁴⁶ The details of the corpora including all the datasets involved in training and testing the models for each language are shown in Table 17.

Kurdish (including Sorani)

The amount of Sorani-specific parallel Sorani-English data available online is rather small, therefore we opted to also include general Kurdish-English parallel corpora (available under kur macrolanguage ISO 639-3 code) in our training data. The resulting model is in essence a multi-dialect and multi-script model for translating from Sorani (ckb), Kurmanji (kmr), or Northern Kurdish (that uses Latin script), and possibly other Kurdish dialects into English, but we make sure that we evaluate the model using test sentence pairs that include a substantial proportion of source sentences in Perso-Arabic script.

The training set includes the Sorani (ckb) part of XLENT (El-Kishky et al., 2021), TICO-19 (Anastasopoulos et al., 2020),⁴⁷ and Wikimedia (Tiedemann, 2012),⁴⁸ as well as the general Kurdish (kur) training portions of OPUS-100 (Zhang, Williams, Titov, and Sennrich, 2020),⁴⁹ the TANZIL corpus of religious texts,⁵⁰ and the TATOEBE dataset.⁵¹ The development set consists of the development portion of OPUS-100 for general Kurdish. The test set is a general Kurdish test set of OPUS-100 dataset. In this set there are 329 source (Kurdish) sentences out of 2,000 which are in Perso-Arabic script.

Sindhi

The Sindhi training data includes the Sindhi-English parallel data from the following datasets: XLENT, Wikimedia, TANZIL, CCMATRIX (Fan et al., 2021; Schwenk et al., 2021),⁵² and QED (Abdelali, Guzman, Sajjad, and Vogel, 2014).⁵³ For the development set we selected the Sindhi-English localization strings from the Ubuntu project.⁵⁴ The test set consists of 1,000 sentence pairs withheld from the OPUS-100 training set.

46. <https://opus.nlpl.eu/>

47. <https://opus.nlpl.eu/tico-19.php>

48. <https://opus.nlpl.eu/wikimedia.php>

49. <https://opus.nlpl.eu/opus-100.php>

50. <https://opus.nlpl.eu/Tanzil.php> (<https://tanzil.net/>)

51. <https://opus.nlpl.eu/Tatoeba.php> (<https://tatoeba.org/en/>)

52. <https://opus.nlpl.eu/CCMatrix.php>

53. <https://opus.nlpl.eu/QED.php>

54. <https://opus.nlpl.eu/Ubuntu.php>

Urdu

In the training data we include the Urdu-English sentence pairs from the ANUVAAD corpus (Anuvaad, 2022), the parallel South Asian corpora from the JOSHUA statistical machine translation (SMT) toolkit (Post, Callison-Burch, and Osborne, 2012),⁵⁵ and the training set from OPUS-100. The development and test sets consist of the respective development and test Urdu-English partitions of the OPUS-100 dataset.

Uyghur

The training data includes the Uyghur-English pairs from the following corpora: XLENT, TANZIL, TATOEB, TED (Reimers and Gurevych, 2020),⁵⁶ and the training partition of OPUS-100. The development and test sets consist of the respective development and test Uyghur-English partitions of the OPUS-100 dataset.

Multilingual Configuration

In addition to constructing individual monolingual translation models we also experiment with a single multilingual model that provides many-to-English translation. Rather than using all the available data we constructed a corpus that is balanced in terms of per-language parallel sentence pairs: all of the training data is selected for Uyghur, which is our smallest dataset (see Table 17), and for the rest of the languages we selected the first 200,000 sentence pairs from the respective training sets. The resulting multilingual training set thus constructed consists of 776,179 sentence pairs. The test set for multilingual configuration consists of 7,000 sentence pairs that correspond to the whole test sets for the respective languages.

5.2.2. *Models*

Modern neural machine translation (NMT) models are an instance of neural sequence-to-sequence models, which have achieved impressive results in recent years (Stahlberg, 2020). In our experiments we use a variant of recurrent neural network (RNN) encoder-decoder bipartite architecture equipped with attention mechanism (Bahdanau, Cho, and Bengio, 2015; Mnih, Heess, Graves, and Kavukcuoglu, 2014), where instead of RNN units, long short-term memory (LSTM) cells are used, which allows the network to learn the long sequences more efficiently (Hochreiter and Schmidhuber, 1997). The particular attention mechanism we use in the decoder is described in Luong, Pham, and Manning (2015).

55. <https://github.com/joshua-decoder/indian-parallel-corpora>

56. <https://opus.nlpl.eu/TED2020.php>

We use two different model configurations. For languages with larger amount of data (Sindhi and Urdu, as shown in Table 17), the encoder component is bidirectional (Schuster and Paliwal, 1997), consisting of four stacked layers of 256 LSTM units each, while the decoder memory consists of 512 units. The configuration used for languages with smaller amounts of data, Kurdish and Uyghur, is mostly identical, but the encoder has two unidirectional LSTM layers. Both models correspond to vanilla configurations (NMTMediumV1 and NMTSmallV1) provided by the OPENNMT-TF library (Klein et al., 2017) implemented in the TensorFlow framework (Abadi et al., 2016). Our models are word-based, with the 50,000 most frequent words used for source and target embedding vocabularies. The tokenization is performed in *aggressive* mode provided by the default OPENNMT tokenizer. The parallel sentence pairs where either the source or the target sentence is longer than 100 words are dropped from the training. Overall, the larger models NMTMediumV1 have approximately 92M model parameters, while the smaller models NMTSmallV1 have approximately 62M parameters. We used default hyperparameters provided by the OPENNMT configurations, apart from the training batch size which we set to 64 examples.

For the multilingual experiment, the size of the balanced dataset described in §5.2.1 roughly corresponds to the size of our Urdu corpus. Hence, similar to Urdu and Sindhi, we have chosen the NMTMediumV1 configuration for our multilingual many-to-English model with the same hyperparameters as for the monolingual configurations.

5.2.3. Results and Discussion

For each language two native language-to-English models were trained from unnormalized and normalized text for that language, respectively.⁵⁷ The details of the language-specific text partitions are provided in §5.2.1. Perso-Arabic script normalization was applied to the native language side of training, development and testing portions of the data, with the English side kept unchanged. Each model was trained for 8 epochs and at the end of each epoch model’s performance was evaluated on the test set using the three MT metrics, each using default parameters such as casing and smoothing, provided by the SACRE-BLEU toolkit (Post, 2018):⁵⁸ the BiLingual Evaluation Understudy, or BLEU score (Papineni, Roukos, Ward, and Zhu, 2002), the Character n -gram F-score, or CHRF2 (Popović, 2016), and Translation Edit Rate,

57. While we do not provide the statistics for the normalized NMT data in terms of number of training set lines, tokens and types changed by the normalization, we hypothesize that these ratios would be similar to the ones computed for statistical language modeling experiments using Wikipedia data presented in Table 12 in §5.1.

58. <https://github.com/mjpost/sacrebleu>

TABLE 18. Relative difference (%) between the performance of normalized and unnormalized models

(a) Kurdish				(b) Sindhi			
Epochs	Δ BLEU	Δ CHRF2	Δ TER	Epochs	Δ BLEU	Δ CHRF2	Δ TER
1	-3.216	30.460	0.665	1	11.239	2.264	4.059
2	-10.719	-5.340	-0.883	2	6.358	2.351	-3.655
3	2.249	8.010	-3.095	3	3.536	0.028	-2.644
4	30.132	14.786	-2.996	4	7.950	3.158	-2.161
5	28.440	17.847	1.985	5	2.024	1.542	4.414
6	20.165	16.023	-2.022	6	1.075	0.652	-0.503
7	17.866	12.143	-6.269	7	3.384	3.173	1.117
8	31.357	18.202	-4.183	8	1.254	-0.298	-2.795
μ	14.534	14.015	-2.099	μ	4.603	1.609	-0.271

(c) Urdu				(d) Uyghur			
Epochs	Δ BLEU	Δ CHRF2	Δ TER	Epochs	Δ BLEU	Δ CHRF2	Δ TER
1	-2.063	-2.960	-2.657	1	-5.998	-1.583	0.130
2	10.752	3.428	-1.981	2	5.920	1.378	2.780
3	6.686	1.269	-2.686	3	2.661	0.960	0.128
4	8.741	4.109	-0.764	4	6.525	7.454	-5.967
5	4.781	1.423	0.121	5	1.874	0.046	1.550
6	3.657	0.610	-1.366	6	-2.120	-3.592	-2.101
7	4.634	1.606	-2.558	7	-0.345	1.451	0.748
8	3.367	1.726	-0.441	8	0.600	0.348	1.224
μ	5.069	1.401	-1.541	μ	1.140	0.808	-0.188

or TER (Snover et al., 2006). Higher BLEU and CHRF2 scores indicate that the hypotheses better match the reference translations, whereas for TER lower scores indicate a better match.

The relative differences (in %) computed between the scores for the models build on normalized and unnormalized text for each language are shown in Table 18, with the positive values of Δ BLEU and Δ CHRF2, and the negative values of Δ TER signifying relative improvement in performance of the normalized model over the unnormalized one at each training epoch. The last two highlighted rows in each table correspond to relative performance differences at the last epoch and the mean per-epoch difference μ . The improvements are highlighted in green and the degradation in red.⁵⁹

As can be seen from Table 18, the biggest gains over the baseline are obtained for the normalized Kurdish model over all the three MT metrics for both the last training epoch as well as the average per-epoch relative difference. The normalized Urdu model also displays improvements across the board. For Sindhi, there is a relative degradation of 0.298%

59. The absolute raw scores are also provided in Table 19.

TABLE 19. Paired significance tests for the monolingual models obtained after the final training epoch: Paired Bootstrap Resampling (PBS) and Paired Approximate Randomization (PAR)

Systems		BLEU ($\mu \pm 95\%$ CI)	PBS			PAR		
			CHRF2 ($\mu \pm 95\%$ CI)	TER ($\mu \pm 95\%$ CI)	BLEU	CHRF2	TER	
Kurdish	\mathcal{B}	10.740 (10.731 \pm 1.082)	27.187 (27.193 \pm 1.068)	83.047 (83.070 \pm 1.442)	10.740	27.187	83.047	
	\mathcal{N}	15.646 (15.617 \pm 1.269)	33.237 (33.220 \pm 1.210)	79.712 (79.742 \pm 1.769)	15.646	33.237	79.712	
	f	0.001	0.001	0.001	0.000	0.000	0.000	
Sindhi	\mathcal{B}	15.362 (15.367 \pm 0.675)	39.223 (39.228 \pm 0.707)	74.960 (74.967 \pm 1.034)	15.362	39.223	74.960	
	\mathcal{N}	15.557 (15.572 \pm 0.717)	39.106 (39.116 \pm 0.719)	72.921 (72.923 \pm 0.991)	15.557	39.106	72.921	
	f	0.148	0.195	0.001	0.392	0.575	0.000	
Urdu	\mathcal{B}	13.387 (13.393 \pm 0.702)	33.559 (33.562 \pm 0.771)	77.403 (77.402 \pm 1.002)	13.387	33.559	77.403	
	\mathcal{N}	13.854 (13.854 \pm 0.709)	34.148 (34.151 \pm 0.751)	77.064 (77.061 \pm 1.066)	13.854	34.148	77.064	
	f	0.043	0.005	0.167	0.091	0.010	0.442	
Uyghur	\mathcal{B}	9.982 (9.938 \pm 0.588)	29.118 (29.117 \pm 0.636)	87.236 (87.207 \pm 1.678)	9.982	29.118	87.236	
	\mathcal{N}	10.043 (10.017 \pm 0.581)	29.220 (29.225 \pm 0.628)	88.317 (88.352 \pm 1.413)	10.043	29.220	88.317	
	f	0.315	0.225	0.105	0.825	0.668	0.277	

in CHRF2 over the unnormalized baseline at the last training epoch, while for Uyghur there is a larger final-epoch degradation of 1.224% in TER. Apart from these two cases however, overall the mean per-epoch and final-epoch relative differences indicate potential improvements, although in the case of Uyghur these are small.

In order to ascertain whether the above relative differences are statistically significant we performed paired significance testing of the final-epoch systems using two algorithms provided by SACREBLEU: the Paired Bootstrap Resampling (Koehn, 2004) and the Paired Approximate Randomization (Riezler and Maxwell, 2005), denoted PBS and PAR, respectively. For PBS, the default parameter of 1,000 resampling trials was used. For PAR, the default value of 10,000 trials was used for randomization test. The results of both tests are shown in Table 19. For each language two systems are tested: the unnormalized baseline (\mathcal{B}), and the model built from the normalized text (\mathcal{N}). The systems are pairwise compared for sentences from the test set using the three MT metrics described above. The null hypotheses for both tests postulate that both \mathcal{B} and \mathcal{N} translations are generated by the same underlying process. For a given model \mathcal{N} and the baseline \mathcal{B} , the p -value is roughly the probability of the absolute score difference (Δ) or higher occurring due to chance, under the assumption that the null hypothesis is correct. Assuming a significance threshold of 0.05, the null hypothesis can be rejected for p -values < 0.05 , which implies that both systems are different. For PBS, the actual system score, the bootstrap estimated true mean (μ), and the 95% confidence interval (CI) are shown for each metric. For PAR, no true mean or confidence intervals are shown because the algorithm does not perform the resampling.

In Table 19 the statistically significant improvements are highlighted in green, while the cases where the systems appear to be equivalent are

TABLE 20. Relative difference (%) between the performance of two (normalized and unnormalized) many-to-English models

Epochs	Δ BLEU	Δ CHRF2	Δ TER
1	31.456	12.001	5.988
2	35.790	18.781	-4.837
3	28.537	14.952	-2.812
4	25.456	13.398	-5.170
5	22.387	9.781	-4.486
6	14.101	7.564	-6.108
7	15.732	8.336	-4.087
8	10.666	5.702	-6.660
μ	23.016	11.314	-3.521

highlighted in light blue. We note that both small degradations in translation quality of Sindhi and Uyghur for the individual metrics observed in Table 18 turn out to be not statistically significant, as evidenced by the corresponding p -values. For Kurdish, the improvements are statistically significant across the board, while for Sindhi and Urdu the improvements are significant according to at least one MT metric and at least one significance test: both PBS and PAR agree on improvements in TER on Sindhi and in CHRF2 on Urdu (where PBS also indicates significant improvement in BLEU). Interestingly, both tests indicate that normalization has no effect on Uyghur translation quality. We hypothesize that this may be due to several conflating factors. First, since this is the smallest dataset of all the languages in this experiment (see Table 17), there may not be enough data for training the model reliably. Furthermore, potential misalignment between Uyghur and English sentences in the training data may be adversely affecting the quality of the resulting models.

Many-to-English Experiment

The goal of this experiment is to verify the hypothesis that the relative performance of Perso-Arabic script-normalized individual NMT systems, especially Uyghur, is improved by pooling the data from other available languages. To this end we trained a single many-to-English model described in §5.2.1 and §5.2.2. Similar to individual language-to-English experiments, we compare the performance of the NMT model built from normalized text against the baseline model constructed from unnormalized data.

Before proceeding two important points need to be noted. First, because our Perso-Arabic script normalization grammars are language-specific, the normalized version of the multilingual corpus described in §5.2.1 is constructed from the normalized corpora for the respective individual languages. Second, since our balanced multilingual corpus

TABLE 21. Paired significance tests for the multilingual model obtained after the final training epoch: Paired Bootstrap Resampling (PBS) and Paired Approximate Randomization (PAR)

Systems		PBS			PAR		
		BLEU ($\mu \pm 95\%$ CI)	CHRF2 ($\mu \pm 95\%$ CI)	TER ($\mu \pm 95\%$ CI)	BLEU	CHRF2	TER
Kurdish	\mathcal{B}^m	12.968 (12.891 \pm 1.471)	29.403 (29.412 \pm 1.144)	83.296 (83.300 \pm 2.881)	12.968	29.402	83.296
	\mathcal{N}^m	18.496 (18.505 \pm 1.511)	34.773 (34.781 \pm 1.334)	73.373 (73.371 \pm 1.820)	18.496	34.773	73.373
	p	0.001	0.001	0.001	0.000	0.000	0.000
Sindhi	\mathcal{B}^m	14.602 (14.586 \pm 0.690)	37.889 (37.901 \pm 0.649)	76.918 (76.920 \pm 1.105)	14.602	37.889	76.918
	\mathcal{N}^m	15.715 (15.727 \pm 0.742)	39.410 (39.425 \pm 0.675)	72.889 (72.895 \pm 1.011)	15.715	39.410	72.890
	p	0.001	0.001	0.001	0.000	0.000	0.000
Urdu	\mathcal{B}^m	11.226 (11.211 \pm 0.621)	30.516 (30.526 \pm 0.677)	84.246 (84.236 \pm 1.651)	11.226	30.516	84.246
	\mathcal{N}^m	12.255 (12.247 \pm 0.658)	32.057 (32.057 \pm 0.721)	79.169 (79.185 \pm 1.065)	12.255	32.057	79.169
	p	0.001	0.001	0.001	0.000	0.000	0.000
Uyghur	\mathcal{B}^m	15.346 (15.361 \pm 0.820)	35.412 (35.419 \pm 0.788)	75.741 (75.731 \pm 1.206)	15.346	35.412	75.741
	\mathcal{N}^m	17.031 (17.050 \pm 0.881)	37.377 (37.388 \pm 0.877)	71.677 (71.656 \pm 1.182)	17.031	37.377	71.677
	p	0.001	0.001	0.001	0.000	0.000	0.000

has roughly the same number of sentence pairs for each language, the amount of within-language data available for the “bigger” languages (Urdu and Sindhi) is significantly smaller in this experiment. Thus, when one compares NMT scores for these languages between a many-to-English system on one hand, and a particular monolingual language-to-English system on the other, the many-to-English scores may be worse. This is not an issue because, as mentioned above, the goal of this experiment is to investigate the relative improvements over the unnormalized baseline, rather than constructing an NMT system with the best possible absolute score.

The relative differences (in %) computed between the scores for the many-to-English models build from normalized and unnormalized text are shown in Table 20. It is worth noting that unlike the monolingual scores shown in Table 18 these scores are computed using the combined test data consisting of 7,000 sentence pairs from all the individual languages described in §5.2.1. As can be seen from Table 20, the normalized many-to-English model shows consistent improvements in all the metrics over the unnormalized baseline for all the epochs with the exception of 5.988% degradation in TER for the initial epoch.

We also performed PBS and PAR paired statistical significance tests for the many-to-English configuration comparing the performance of the multilingual normalized model (denoted \mathcal{N}^m) against its unnormalized counterpart (\mathcal{B}^m) on the test data for individual languages. The results of both tests are shown in Table 21 for each language, with the statistically significant improvements in individual metrics marked in green. Compared to significance tests for the monolingual systems in Table 19, the multilingual tests show more robust improvements across all languages and metrics. In particular, with respect to Uyghur these results confirm our hypothesis above that the original dataset is

too small to reliably measure the effects of script normalization. This is rectified by using data from other languages. Additional supporting evidence comes from comparing between the absolute values of all the metrics for normalized and unnormalized monolingual and multilingual models for Uyghur and Kurdish shown in Table 19 and Table 21: multilingual configurations have higher absolute scores for these languages.

6. Conclusions

This paper provided a brief overview of various adaptations of the Perso-Arabic script for eight languages from diverse language families and the issues that result from representing these adaptations digitally. The particular emphasis of this study was on the visual ambiguities between Perso-Arabic characters represented in Unicode. We argued that the computational methods for visual disambiguation need to go beyond the standard language-agnostic techniques provided by the Unicode standard and take into account the specifics of a local writing system as well as multiple confounding factors that affect the patterns of its use. We presented two types of writing system-specific normalization methods. Similar to the standard Unicode normalization techniques, *visual* normalization preserves the visual invariance of the characters, while providing significantly broader coverage of normalization cases peculiar to the orthography in question. The second type is *reading* normalization, which provides character transformations that violate visual invariance (e.g., by modifying the number of *ijām* dots on the base shape of a character), yet are required to make the input conform to the local orthography. The distinction between the two types hinges on the visual invariance criterion, which is helpful in deciding when and if to apply either type of technique.⁶⁰ Perso-Arabic script normalization techniques are crucial for cybersecurity, but the focus of this paper is on their application to natural language processing. We performed experiments in statistical language modeling and neural machine translation that demonstrated the positive impact of script normalization on the performance of the resulting models.

This study describes work that is still in early stages. While there is a wealth of literature on the eight languages described in this paper and some additional languages currently covered by our methods, the majority of Perso-Arabic writing systems are used for lower-resource languages that are either scarcely documented or have very little online data, which is needed to provide evidence for required normalization.

60. In some applications it may be necessary to preserve the visual fidelity of the input, hence the application of reading normalization may not be desirable.

Furthermore, significant research towards a formal description of Perso-Arabic script typology is still required.

Acknowledgements

The authors would like to thank Christo Kirov for many useful comments on an earlier draft of this paper, as well as Lawrence Wolf-Sonkin, Işın Demirşahin and Anna Katanova for informative feedback and assistance with the various stages of this project. We also thank Aso Mahmudi for his invaluable feedback on the modern orthography of Central Kurdish.

A. Letter Inventories for Individual Languages

The letter repertoire for the eight languages investigated in this paper—South Azerbaijani (azb), Sorani Kurdish (ckb), Kashmiri (ks), Malay (ms), Western Punjabi (pnb), Sindhi (sd), Uyghur (ug) and Urdu (ur)—is shown in Table 22 below. Overall we identified 118 characters including letters and various diacritics. The table shows, for each character its corresponding Unicode code point, Unicode name and the languages that use it (indicated by the checkmark).

TABLE 22. Letter inventories for individual languages

Char.	Codepoint	Character Name	Language Tags							
			azb	ckb	ks	ms	pnb	sd	ug	ur
<ي>	U+0620	Kashmiri Yeh			✓					
<ء>	U+0621	Hamza				✓	✓	✓		✓
<ا>	U+0622	Alef with Madda Above	✓		✓	✓	✓	✓		✓
<آ>	U+0623	Alef with Hamza Above			✓	✓	✓	✓		✓
<ؤ>	U+0624	Waw with Hamza Above	✓		✓	✓	✓	✓		✓
<إ>	U+0625	Alef with Hamza Below			✓	✓	✓			✓
<آ>	U+0626	Yeh with Hamza Above	✓	✓			✓	✓	✓	✓
<ا>	U+0627	Alef	✓	✓	✓	✓	✓	✓	✓	✓
<ب>	U+0628	Beh	✓	✓	✓	✓	✓	✓	✓	✓
<ة>	U+0629	Teh Marbuta				✓	✓		✓	✓
<ت>	U+062A	Teh	✓	✓	✓	✓	✓	✓	✓	✓
<ٹ>	U+062B	Theh	✓		✓	✓	✓	✓		✓
<ج>	U+062C	Jeem	✓	✓	✓	✓	✓	✓	✓	✓
<ح>	U+062D	Hah	✓	✓	✓	✓	✓	✓	✓	✓
<خ>	U+062E	Khah	✓	✓	✓	✓	✓	✓	✓	✓

Continued on next page

TABLE 22—Continued from previous page

Char.	Codepoint	Character Name	Language Tags							
			az	ckb	ks	ms	pa	sd	ug	ur
<د>	U+062F	Dal	✓	✓	✓	✓	✓	✓	✓	✓
<ذ>	U+0630	Thal	✓		✓	✓	✓	✓	✓	✓
<ر>	U+0631	Reh	✓	✓	✓	✓	✓	✓	✓	✓
<ز>	U+0632	Zain	✓	✓	✓	✓	✓	✓	✓	✓
<س>	U+0633	Seen	✓	✓	✓	✓	✓	✓	✓	✓
<ش>	U+0634	Sheen	✓	✓	✓	✓	✓	✓	✓	✓
<ص>	U+0635	Sad	✓		✓	✓	✓	✓		✓
<ض>	U+0636	Dad	✓		✓	✓	✓	✓		✓
<ط>	U+0637	Tah	✓		✓	✓	✓	✓		✓
<ظ>	U+0638	Zah	✓		✓	✓	✓	✓		✓
<ع>	U+0639	Ain		✓	✓	✓	✓	✓	✓	✓
<غ>	U+063A	Ghain	✓	✓	✓	✓	✓	✓	✓	✓
<ی>	U+063D	Farsi Yeh with Inverted V	✓							
<ف>	U+0641	Feh	✓	✓	✓	✓	✓	✓	✓	✓
<ق>	U+0642	Qaf	✓	✓	✓	✓	✓	✓	✓	✓
<ك>	U+0643	Kaf							✓	✓
<ل>	U+0644	Lam	✓	✓	✓	✓	✓	✓	✓	✓
<م>	U+0645	Meem	✓	✓	✓	✓	✓	✓	✓	✓
<ن>	U+0646	Noon	✓	✓	✓	✓	✓	✓	✓	✓
<ه>	U+0647	Heh	✓	✓		✓	✓			✓
<و>	U+0648	Waw	✓	✓	✓	✓	✓	✓	✓	✓
<ى>	U+0649	Alef Maksura		✓		✓	✓		✓	✓
<ي>	U+064A	Yeh				✓		✓	✓	
<َ>	U+064B	Fathatan					✓			✓
<ِ>	U+064C	Dammatan					✓			✓
<ُ>	U+064D	Kasratan					✓			✓
<ا>	U+064E	Fatha	✓		✓	✓	✓	✓	✓	✓
<آ>	U+064F	Damma	✓		✓	✓	✓	✓	✓	✓
<إ>	U+0650	Kasra	✓		✓	✓	✓	✓	✓	✓
<ع>	U+0651	Shadda	✓		✓	✓	✓	✓		✓
<ـ>	U+0652	Sukun	✓		✓	✓	✓			✓
<ْ>	U+0653	Maddah Above	✓		✓	✓	✓	✓		✓
<َٓ>	U+0654	Hamza Above			✓		✓	✓	✓	✓
<ِٓ>	U+0655	Hamza Below			✓		✓	✓		✓
<ُٓ>	U+0656	Subscript Alef			✓		✓			
<آ>	U+0657	Inverted Damma			✓		✓			✓
<ٓ>	U+065A	Vowel Sign Small V Above			✓					
<ٓٓ>	U+065F	Wavy Hamza Below			✓					
<ٓٔ>	U+0670	Superscript Alef					✓			✓
<آٔ>	U+0671	Alef Wasla	✓							
<آٔٓ>	U+0672	Alef with Wavy Hamza Above			✓					
<آٔٔ>	U+0673	Alef with Wavy Hamza Below			✓					
<ت>	U+0679	Tteh			✓		✓			✓
<تٓ>	U+067A	Tteheh						✓		

Continued on next page

TABLE 22—Continued from previous page

Char.	Codepoint	Character Name	Language Tags							
			az	ckb	ks	ms	pa	sd	ug	ur
<ٲ>	U+067B	Beeh						✓		
<ٲ̣>	U+067D	Teh with Three Dots Above Downwards						✓		
<ٲ̣>	U+067E	Peh	✓	✓	✓		✓	✓	✓	✓
<ٲ̣̣>	U+067F	Teheh						✓		
<ٲ̣̣̣>	U+0680	Beheh						✓		
<ٲ̣̣̣̣>	U+0683	Nyeh						✓		
<ٲ̣̣̣̣̣>	U+0684	Dyeh						✓		
<ٲ̣̣̣̣̣̣>	U+0685	Hah with Three Dots Above							✓	
<ٲ̣̣̣̣̣̣̣>	U+0686	Tcheh	✓	✓	✓	✓	✓	✓	✓	✓
<ٲ̣̣̣̣̣̣̣̣>	U+0687	Tcheheh						✓		
<ٲ̣̣̣̣̣̣̣̣̣>	U+0688	Ddal			✓		✓			✓
<ٲ̣̣̣̣̣̣̣̣̣̣>	U+068A	Dal with Dot Below						✓		
<ٲ̣̣̣̣̣̣̣̣̣̣̣>	U+068C	Dahal						✓		
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣>	U+068D	Ddahal						✓		
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+068F	Dal with Three Dots Above Downwards						✓		
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+0691	Rreh			✓		✓			✓
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+0695	Reh with Small V Below		✓						
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+0698	Jeh	✓	✓	✓		✓		✓	✓
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+0699	Reh with Four Dots Above						✓		
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06A0	Ain with Three Dots Above				✓				
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06A4	Veh		✓		✓				
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06A6	Peheh						✓		
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06A9	Keheh	✓	✓	✓	✓	✓	✓		✓
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06AA	Swash Kaf						✓		
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06AD	Ng							✓	
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06AF	Gaf	✓	✓	✓		✓	✓	✓	✓
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06B1	Ngoeh						✓		
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06B3	Gueh						✓		
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06B4	Gaf with Three Dots Above	✓							
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06B5	Lam with Small V		✓						
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06BA	Noon Ghunna			✓		✓			✓
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06BB	Rnoon						✓		
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06BD	Noon with Three Dots Above				✓				
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06BE	Heh Doachashmee	✓	✓			✓	✓	✓	✓
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06C0	Heh with Yeh Above	✓							
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06C1	Heh Goal			✓		✓			✓
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06C2	Heh Goal with Hamza Above			✓		✓			✓
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06C3	Teh Marbuta Goal					✓			✓
<ٲ̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣̣>	U+06C4	Waw with Ring			✓					

Continued on next page

TABLE 22—Continued from previous page

Char.	Codepoint	Character Name	Language Tags							
			az	ckb	ks	ms	pa	sd	ug	ur
<ۆ>	U+06C6	Oe	✓	✓	✓				✓	
<ۇ>	U+06C7	U	✓						✓	
<ۈ>	U+06C8	Yu							✓	
<ۋ>	U+06CA	Waw with Two Dots Above		✓						
<ۋ>	U+06CB	Ve							✓	
<ۏ>	U+06CC	Farsi Yeh	✓	✓	✓		✓			✓
<ۏ>	U+06CE	Yeh with Small V	✓	✓						
<ۋ>	U+06CF	Waw with Dot Above				✓				
<ۏ>	U+06D0	E							✓	
<ۏ>	U+06D2	Yeh Barree					✓			✓
<ۏ>	U+06D3	Yeh Barree with Hamza Above					✓			✓
<ۏ>	U+06D5	Ae		✓					✓	
<ۏ>	U+06FD	Sign Sindhi Ampersand						✓		
<ۏ>	U+06FE	Sign Sindhi Postposition Men						✓		
<ۏ>	U+0762	Keheh with Dot Above				✓				
<ۏ>	U+0762	Keheh with Dot Above				✓				
<ۏ>	U+0763	Keheh with Three Dots Above				✓				
<ۏ>	U+0768	Noon with Small Tah Above					✓			
<ۏ>	U+076C	Reh with Hamza Above			✓					
<ۏ>	U+08C7	Lam with Small Tah Above					✓			

B. Language Model Experiments

B.1. Character Language Models

Full character language model results are shown here for Kashmiri (Table 23), Kurdish Sorani (Table 24), Malay (Table 25), Punjabi Shahmukhi (Table 26), Sindhi (Table 27), South Azerbaijani (Table 28), Uyghur (Table 29), and Urdu (Table 30).

Assuming the significance level of $\alpha = 0.05$, the results for Kashmiri in Table 23, which is the smallest dataset, indicate that cross-entropy improvements for n -gram orders 4, 9, 10 are statistically significant. For the 8-grams, the MW and BW tests indicate borderline significance

TABLE 23. Significance tests for Kashmiri character n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
3	-0.003	0.08	-0.015	0.009	-0.458	0.647	5242.0	0.555	-0.588	0.557
4	-0.021	0.71	-0.036	-0.006	-2.722	0.007	5946.0	0.021	-2.356	0.019
5	-0.006	0.21	-0.024	0.013	-0.607	0.544	5247.0	0.547	-0.601	0.549
6	-0.006	0.22	-0.023	0.012	-0.652	0.515	5327.0	0.425	-0.793	0.429
7	-0.005	0.2	-0.024	0.013	-0.578	0.564	5276.0	0.501	-0.671	0.503
8	-0.016	0.63	-0.034	0.002	-1.791	0.075	5802.0	0.05	-1.977	0.05
9	-0.022	0.85	-0.04	-0.004	-2.468	0.014	6007.0	0.014	-2.51	0.013
10	-0.028	1.07	-0.046	-0.009	-2.96	0.003	6065.0	0.009	-2.676	0.008

TABLE 24. Significance tests for Kurdish (Sorani) character n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
3	-0.01	0.36	-0.011	-0.009	-22.395	0.0	9916.0	0.0	-115.778	0.0
4	-0.003	0.16	-0.004	-0.002	-6.687	0.0	7535.0	0.0	-7.394	0.0
5	-0.004	0.24	-0.005	-0.003	-6.288	0.0	7369.0	0.0	-6.724	0.0
6	-0.003	0.2	-0.004	-0.002	-4.633	0.0	6718.0	0.0	-4.547	0.0
7	-0.004	0.28	-0.005	-0.002	-5.768	0.0	7270.0	0.0	-6.294	0.0
8	-0.006	0.44	-0.007	-0.004	-8.596	0.0	8040.0	0.0	-10.003	0.0
9	-0.005	0.43	-0.006	-0.004	-8.741	0.0	8006.0	0.0	-9.734	0.0
10	-0.005	0.39	-0.006	-0.003	-7.075	0.0	7623.0	0.0	-7.86	0.0

TABLE 25. Significance tests for Malay character n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
3	-0.065	1.818	-0.066	-0.064	-194.951	0.0	10 000.0	0.0	$-\infty$	0.0
4	-0.062	2.036	-0.063	-0.061	-143.242	0.0	10 000.0	0.0	$-\infty$	0.0
5	-0.06	2.232	-0.061	-0.059	-135.144	0.0	10 000.0	0.0	$-\infty$	0.0
6	-0.064	2.526	-0.065	-0.063	-132.604	0.0	10 000.0	0.0	$-\infty$	0.0
7	-0.067	2.736	-0.068	-0.066	-135.678	0.0	10 000.0	0.0	$-\infty$	0.0
8	-0.07	2.885	-0.071	-0.069	-126.796	0.0	10 000.0	0.0	$-\infty$	0.0
9	-0.07	2.931	-0.071	-0.069	-116.01	0.0	10 000.0	0.0	$-\infty$	0.0
10	-0.07	2.922	-0.071	-0.068	-103.192	0.0	10 000.0	0.0	$-\infty$	0.0

TABLE 26. Significance tests for Punjabi (Shahmukhi) character n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
3	-0.011	0.32	-0.012	-0.01	-34.213	0.0	10 000.0	0.0	$-\infty$	0.0
4	-0.013	0.44	-0.013	-0.012	-35.959	0.0	10 000.0	0.0	$-\infty$	0.0
5	-0.01	0.39	-0.011	-0.009	-24.399	0.0	9941.0	0.0	-167.411	0.0
6	-0.008	0.34	-0.009	-0.007	-16.158	0.0	9667.0	0.0	-29.955	0.0
7	-0.007	0.33	-0.008	-0.006	-13.856	0.0	9429.0	0.0	-21.268	0.0
8	-0.006	0.3	-0.008	-0.004	-6.491	0.0	8928.0	0.0	-14.653	0.0
9	-0.007	0.36	-0.009	-0.006	-7.799	0.0	9061.0	0.0	-16.395	0.0
10	-0.007	0.34	-0.009	-0.005	-6.11	0.0	8609.0	0.0	-12.104	0.0

TABLE 27. Significance tests for Sindhi character n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
3	0.001	-0.03	-0.002	0.004	0.71	0.479	4649.0	0.392	0.853	0.395
4	0.011	-0.28	0.005	0.016	3.628	0.0	3461.0	0.0	3.935	0.0
5	-0.018	0.47	-0.028	-0.008	-3.451	0.001	6718.0	0.0	-4.51	0.0
6	-0.028	0.79	-0.039	-0.018	-5.183	0.0	7477.0	0.0	-7.109	0.0
7	-0.016	0.46	-0.025	-0.006	-3.309	0.001	6714.0	0.0	-4.448	0.0
8	-0.024	0.71	-0.033	-0.014	-4.775	0.0	6864.0	0.0	-4.977	0.0
9	-0.015	0.45	-0.026	-0.004	-2.715	0.007	6119.0	0.006	-2.814	0.005
10	-0.014	0.42	-0.026	-0.003	-2.443	0.015	5889.0	0.03	-2.212	0.028

TABLE 28. Significance tests for South Azerbaijani character n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
3	-0.012	0.418	-0.013	-0.011	-24.371	0.0	9953.0	0.0	-192.439	0.0
4	-0.004	0.186	-0.005	-0.003	-6.649	0.0	7482.0	0.0	-7.222	0.0
5	-0.002	0.131	-0.003	-0.001	-3.133	0.002	6137.0	0.005	-2.866	0.005
6	-0.001	0.093	-0.003	0.0	-1.993	0.048	5837.0	0.041	-2.072	0.04
7	-0.002	0.14	-0.003	0.0	-2.526	0.012	5895.0	0.029	-2.227	0.027
8	-0.002	0.143	-0.003	0.0	-2.52	0.013	5997.0	0.015	-2.489	0.014
9	-0.002	0.13	-0.003	0.0	-2.348	0.02	5945.0	0.021	-2.352	0.02
10	-0.002	0.171	-0.004	-0.001	-2.889	0.004	6238.0	0.002	-3.107	0.002

TABLE 29. Significance tests for Uyghur character n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
3	-0.002	0.074	-0.003	-0.001	-4.42	0.0	6636.0	0.0	-4.273	0.0
4	-0.001	0.051	-0.002	0.0	-2.257	0.025	5801.0	0.05	-1.981	0.049
5	-0.004	0.203	-0.005	-0.003	-7.131	0.0	7690.0	0.0	-7.964	0.0
6	-0.004	0.219	-0.005	-0.003	-6.315	0.0	7394.0	0.0	-6.829	0.0
7	-0.003	0.164	-0.004	-0.002	-4.348	0.0	6661.0	0.0	-4.353	0.0
8	-0.004	0.225	-0.005	-0.002	-5.126	0.0	6907.0	0.0	-5.133	0.0
9	-0.004	0.24	-0.006	-0.002	-4.549	0.0	6699.0	0.0	-4.476	0.0
10	-0.005	0.295	-0.006	-0.003	-6.446	0.0	7460.0	0.0	-6.977	0.0

TABLE 30. Significance tests for Urdu character n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
3	-0.003	0.11	-0.004	-0.001	-4.372	0.0	6646.0	0.0	-4.306	0.0
4	-0.005	0.22	-0.005	-0.004	-17.256	0.0	9575.0	0.0	-34.762	0.0
5	-0.004	0.2	-0.004	-0.003	-9.06	0.0	8813.0	0.0	-15.737	0.0
6	-0.004	0.26	-0.006	-0.003	-5.352	0.0	8363.0	0.0	-11.571	0.0
7	-0.004	0.24	-0.005	-0.002	-4.601	0.0	8246.0	0.0	-10.74	0.0
8	-0.002	0.14	-0.004	0.0	-2.35	0.02	8220.0	0.0	-10.431	0.0
9	-0.002	0.13	-0.005	0.001	-1.24	0.217	7896.0	0.0	-8.207	0.0
10	-0.005	0.37	-0.008	-0.002	-3.422	0.001	8344.0	0.0	-10.745	0.0

disagreeing with the WS test. The discrepancy observed between lower orders is probably due to overfitting as the dataset is tiny. Punjabi (Shahmukhi) and Kurdish (Sorani) are the second and third biggest datasets, respectively, and the results in Table 26 and Table 24 indicate statistically significant improvements across the board, with all the three tests agreeing with each other. Sindhi (Table 27) is similar in some respects to Kashmiri in that the low n -gram orders of 3 and 4 are not very reliable, while the results for the rest of the orders indicate significant improvements. While overfitting may play a certain role, upon informal inspection it appears that, similar to Kashmiri, the Sindhi dataset is quite noisy, even after filtering. Finally, the results for Urdu in Table 30 indicate that, similar to Kurdish (Sorani), Punjabi (Shahmukhi), South Azerbaijani and Uyghur the improvements are statistically significant across the board. The best results are obtained for Malay (Table 25) with up to 2.9% improvement in character entropy.

B.2. Word Language Models

Full word language model results are shown here for Kashmiri (Table 31), Kurdish Sorani (Table 32), Malay (Table 33), Punjabi Shahmukhi (Table 34), Sindhi (Table 35), South Azerbaijani (Table 36), Uyghur (Table 37), and Urdu (Table 38).

As can be seen from Table 31, the hypothesis testing for Kashmiri shows that the null hypothesis is confirmed by all the three algorithms for all the orders n most of the time. This is evident from the tests' p -values (these exceed the significance level $\alpha = 0.05$) as well as the t -test confidence intervals that contain the null hypothesis, which in our case corresponds to the zero difference in means. This is likely the artifact of the models overfitting a very small dataset, even for relatively less scarce bigrams. It is interesting to note that for the rest of the languages the alternative hypotheses for all the models is uniformly confirmed: the small decrease in cross-entropy expressed as bits per word observed for all the languages and configurations is statistically significant.

While the relative improvements for Kurdish Sorani (Table 32), Punjabi Shahmukhi (Table 34), South Azerbaijani (Table 36), Uyghur (Table 37) and Urdu (Table 38) are relatively tiny, possibly due to the relatively small number of modifications compared to the overall size of the datasets, the relative improvements to Sindhi models (Table 35) are over one percent for all the configurations. This may indeed be correlated with the highest number of per-word token modifications for Sindhi among all the languages, denoted R_w in Table 12. The best results are obtained for Malay (Table 33), with up to 3.5% improvement in word entropy.

References

- Aazim, Muzaffar, Kamal Mansour, and Roozbeh Pournader (2009). *Proposal to add two Kashmiri characters and one annotation to the Arabic block*. Tech. rep. L2/09-176. Unicode Consortium.
- Abadi, Martín et al. (2016). "TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems." In: *CoRR* abs/1603.04467.
- Abdelali, Ahmed et al. (2014). "The AMARA Corpus: Building Parallel Language Resources for the Educational Domain." In: *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*. Reykjavik, Iceland: European Language Resources Association (ELRA), pp. 1856–1862.
- Abdullah, Farhanah et al. (2020). "Jawi Script and the Malay Society: Historical Background and Development." In: *International Journal of Management (IJM)* 11.7, 68–78.

TABLE 31. Significance tests for Kashmiri word n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
2	0.034	-0.33	-0.008	0.076	1.595	0.112	4301.0	0.088	1.72	0.087
3	-0.039	0.39	-0.088	0.011	-1.531	0.127	5589.0	0.15	-1.447	0.149
4	-0.022	0.22	-0.067	0.023	-0.957	0.34	5308.0	0.452	-0.751	0.454
5	0.017	-0.17	-0.029	0.063	0.727	0.468	4817.0	0.656	0.443	0.659

TABLE 32. Significance tests for Kurdish (Sorani) word n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
2	-0.031	0.41	-0.038	-0.024	-8.678	0.0	8058.0	0.0	-10.074	0.0
3	-0.034	0.49	-0.04	-0.028	-10.393	0.0	8492.0	0.0	-13.09	0.0
4	-0.034	0.49	-0.041	-0.027	-9.297	0.0	8237.0	0.0	-11.275	0.0
5	-0.035	0.5	-0.042	-0.028	-10.406	0.0	8464.0	0.0	-13.209	0.0

TABLE 33. Significance tests for Malay word n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
2	-0.358	2.935	-0.362	-0.355	-185.4	0.0	10 000.0	0.0	$-\infty$	0.0
3	-0.394	3.319	-0.399	-0.389	-166.053	0.0	10 000.0	0.0	$-\infty$	0.0
4	-0.403	3.411	-0.408	-0.398	-166.567	0.0	10 000.0	0.0	$-\infty$	0.0
5	-0.411	3.479	-0.416	-0.406	-170.479	0.0	10 000.0	0.0	$-\infty$	0.0

TABLE 34. Significance tests for Punjabi (Shahmukhi) word n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
2	-0.015	0.15	-0.018	-0.012	-10.834	0.0	8633.0	0.0	-14.295	0.0
3	-0.017	0.19	-0.021	-0.014	-9.338	0.0	8237.0	0.0	-11.232	0.0
4	-0.02	0.23	-0.025	-0.016	-9.472	0.0	8318.0	0.0	-11.75	0.0
5	-0.02	0.23	-0.024	-0.016	-8.916	0.0	8137.0	0.0	-10.521	0.0

Afzal, Muhammad and Sarmad Hussain (2001). "Urdu computing standards: development of Urdu Zabta Takhti (UZT) 1.01." In: *Proceedings of IEEE International Multi Topic Conference (INMIC): Technology for the 21st Century*. IEEE. Lahore, Pakistan, 216–222.

TABLE 35. Significance tests for Sindhi word n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
2	-0.159	1.06	-0.171	-0.146	-24.103	0.0	9948.0	0.0	-180.99	0.0
3	-0.169	1.21	-0.186	-0.152	-19.861	0.0	9753.0	0.0	-54.902	0.0
4	-0.177	1.26	-0.192	-0.161	-22.847	0.0	9889.0	0.0	-88.059	0.0
5	-0.184	1.32	-0.2	-0.168	-22.57	0.0	9898.0	0.0	-108.596	0.0

TABLE 36. Significance tests for South Azerbaijani word n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
2	-0.031	0.374	-0.038	-0.024	-9.016	0.0	8113.0	0.0	-10.307	0.0
3	-0.027	0.332	-0.034	-0.02	-7.376	0.0	7655.0	0.0	-7.913	0.0
4	-0.031	0.397	-0.038	-0.025	-9.711	0.0	8325.0	0.0	-11.755	0.0
5	-0.028	0.358	-0.035	-0.021	-8.108	0.0	7797.0	0.0	-8.828	0.0

TABLE 37. Significance tests for Uyghur word n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
2	-0.024	0.2	-0.033	-0.015	-5.14	0.0	6961.0	0.0	-5.297	0.0
3	-0.038	0.332	-0.05	-0.027	-6.721	0.0	7537.0	0.0	-7.397	0.0
4	-0.034	0.294	-0.045	-0.022	-5.689	0.0	7177.0	0.0	-6.029	0.0
5	-0.042	0.367	-0.055	-0.029	-6.319	0.0	7381.0	0.0	-6.774	0.0

TABLE 38. Significance tests for Urdu word n -gram language models

n	Δ_μ		WS				MW		BM	
	Δ	%	L	H	t	p	t	p	t	p
2	-0.006	0.08	-0.008	-0.003	-4.923	0.0	6796.0	0.0	-4.804	0.0
3	-0.004	0.06	-0.007	-0.002	-3.507	0.001	6394.0	0.001	-3.531	0.001
4	-0.006	0.1	-0.009	-0.003	-4.2	0.0	6637.0	0.0	-4.285	0.0
5	-0.007	0.1	-0.009	-0.004	-4.927	0.0	6842.0	0.0	-4.914	0.0

Ahmad, Humza and Laszlo Erdodi (2021). “Overview of phishing landscape and homographs in Arabic domain names.” In: *Security and Privacy* 4.4, 1–14.

Ahmadi, Sina (2019). “A rule-based Kurdish text transliteration system.” In: *ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP)* 18.2, 1–8.

- Anastasopoulos, Antonios et al. (2020). "TICO-19: the Translation Initiative for COVID-19." In: *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*. Online: Association for Computational Linguistics.
- Anuvaad, Team (2022). *Project Anuvaad: Parallel Corpus*. Accessed: June, 2022; <https://anuvaaad.org/>.
- Aqtay, Gulayhan (2020). "The New Kazakh Alphabet Based on Latin Script." In: *Language and Society in Kazakhstan: The Kazakh Context*. Ed. by Gulayhan Aqtay and Cem Erdem. Vol. 6. Turkic Studies. Poznań, Poland: Wydawnictwo Naukowe, Adam Mickiewicz University, 23–36.
- Aronson, Howard (1996). "Yiddish." In: *The World's Writing Systems*. Ed. by Peter Daniels and William Bright. Oxford: Oxford University Press. Chap. 61, 735–742.
- Avanzini, Alessandra (2009). "Origin and classification of the Ancient South Arabian languages." In: *Journal of Semitic Studies* 54.1, 205–220.
- Bahdanau, Dmitry, Kyunghyun Cho, and Yoshua Bengio (2015). "Neural machine translation by jointly learning to align and translate." In: *Proceedings of 3rd International Conference on Learning Representations (ICLR)*. San Diego, CA.
- Bashir, Elena (2015). "The Brahui Language: Recovering the Past, Documenting the Present, and Pondering the Future." In: *Proceedings of International Conference on Brahui Language and Culture*. Brahui Academy, Baluchistan, Pakistan. Islamabad, Pakistan, 1–27.
- Bashir, Elena and Thomas J. Connors (2019). *A Descriptive Grammar of Hindko, Panjabi, and Saraiki*. Vol. 4. Mouton-CASL Grammar Series (MCASL). De Gruyter Mouton.
- Bashir, Elena, Sarmad Hussain, and Deborah Anderson (2006). *Proposal for characters for Khowar, Torwali, and Burushaski*. Tech. rep. L2-06/149. Unicode Consortium.
- Bauer, Thomas (1996). "Arabic Writing." In: *The World's Writing Systems*. Ed. by Peter Daniels and William Bright. Oxford: Oxford University Press. Chap. 50, 559–563.
- Bhardwaj, Mangat Rai (2016). *Panjabi: A Comprehensive Grammar*. Routledge Comprehensive Grammars. Routledge.
- Bhatt, Rakesh M. (2015). "Script Choice, Language Loss and the Politics of Anamnesis: Kashmiri in Diaspora." In: *Language, Literacy and Diversity*. Ed. by Christopher Stroud and Mastin Prinsloo. Routledge Critical Studies in Multilingualism. Routledge, 130–147.
- Bosworth, Clifford Edmund (2011). "Arrān." In: *Encyclopædia Iranica*. Ed. by Ehsan Yarshater. Vol. II/5. Online. Leiden, The Netherlands: Brill, 520–522.
- Boudelaa, Sami, Manuel Perea, and Manuel Carreiras (2020). "Matrices of the frequency and similarity of Arabic letters and allographs." In: *Behavior Research Methods* 52.5, 1893–1905.

- Bozarslan, Hamit, Cengiz Gunes, and Veli Yadirgi (2021). "Kurdish Language." In: *The Cambridge History of the Kurds*. Part V. Cambridge, UK: Cambridge University Press, 601–684.
- Braginsky, Vladimir I. (1975). "Some remarks on the structure of the "Sya'ir Perahu" by Hamzah Fansuri." In: *Bijdragen tot de Taal-, Land- en Volkenkunde / Journal of the Humanities and Social Sciences of Southeast Asia and Oceania* 131.4, 407–426.
- Bright, William (1999). "A matter of typology: Alphasyllabaries and abugidas." In: *Written Language & Literacy* 2.1, 45–55.
- Brose, Michael C. (2017). "The Medieval Uyghurs of the 8th through 14th Centuries." In: *Oxford Research Encyclopedia of Asian History*. Online. Oxford University Press, 1–20.
- Brown, Peter F. et al. (1992). "An estimate of an upper bound for the entropy of English." In: *Computational Linguistics* 18.1, 31–40.
- Brunner, Edgar and Ullrich Munzel (2000). "The nonparametric Behrens-Fisher problem: Asymptotic theory and small-sample approximation." In: *Biometrical Journal: Journal of Mathematical Methods in Biosciences* 42.1, 17–25.
- Buljina, Harun (2019). "Empire, Nation, and the Islamic World: Bosnian Muslim Reformists between the Habsburg and Ottoman Empires, 1901–1914." PhD thesis. New York: Columbia University.
- Campbell, George L. (1994). "Kurdish." In: *Concise Compendium of the World's Languages*. Ed. by George L. Campbell. London: Routledge. Chap. 5, 288–292.
- Castilla, Nuria de (2019). "Uses and Written Practices in Aljamiado Manuscripts." In: *Creating Standards: Interactions with Arabic script in 12 manuscript cultures*. Ed. by Dmitry Bondarev, Alessandro Gori, and Lameen Souag. Vol. 16. Studies in Manuscript Cultures. De Gruyter.
- Caswell, Isaac et al. (2020). "Language ID in the Wild: Unexpected Challenges on the Path to a Thousand-Language Web Text Corpus." In: *Proceedings of the 28th International Conference on Computational Linguistics*. Barcelona, Spain (Online): International Committee on Computational Linguistics, 6588–6608.
- Chitralli, Rehmat Aziz Khan (2020a). *Proposal to include Indus Kobistani Language alphabets*. Tech. rep. L2/20-157. Unicode Consortium.
- (2020b). *Proposal to include Kalasba Language alphabets*. Tech. rep. L2/20-091. Unicode Consortium.
- Coluzzi, Paolo (2020). "Jawi, an endangered orthography in the Malaysian linguistic landscape." In: *International Journal of Multilingualism*, 1–17.
- Conneau, Alexis et al. (2020). "Unsupervised Cross-lingual Representation Learning at Scale." In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, 8440–8451.

- Cotterell, Ryan et al. (2018). "Are All Languages Equally Hard to Language-Model?" In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*. New Orleans, Louisiana: Association for Computational Linguistics, 536–541.
- Coulmas, Florian (1999). *The Blackwell Encyclopedia of Writing Systems*. Oxford, UK: John Wiley & Sons.
- Daniels, Peter T. (2013). "The Arabic writing system." In: *The Oxford Handbook of Arabic Linguistics*. Ed. by Jonathan Owens. Oxford, UK: Oxford University Press. Chap. 18, 422–431.
- (2014). "The type and spread of Arabic script." In: *The Arabic Script in Africa*. Ed. by Meikal Mumin and Kees Versteegh. Vol. 71. Studies in Semitic Languages and Linguistics. Brill, 25–39.
- Datta, A. K. (1984). "A Generalized Formal Approach for Description and Analysis of Major Indian Scripts." In: *IETE Journal of Research* 30.6, 155–161.
- Davis, Derek (2015). "Premchand plays chess." In: *Journal of the Royal Asiatic Society* 25.2, 269–300.
- DBP (2006). *Daftar kata bahasa Melayu: Rumi-Sebutan-Jawi [Malay Pronunciation Dictionary (Rumi-Jawi)]*. 5th. Kuala Lumpur, Malaysia: Dewan Bahasa Pustaka (DBP) [Institute of Language and Literature].
- Dillon, Michael (2009). *Xinjiang — China's Muslim Far Northwest*. Durham East Asia Series. London and New York: Routledge.
- Dow, Hugh (1976). "A note on the Sindhi alphabet." In: *Asian Affairs* 7.1, 54–56.
- Dwyer, Arienne M. (2005). *The Xinjiang conflict: Uyghur identity, language policy, and political discourse*. Vol. 15. Policy Studies. Washington, D.C.: East-West Center Washington.
- Elsayed, Yahia and Ahmed Shosha (2018). "Large scale detection of IDN domain name masquerading." In: *Proceedings of 2018 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE. San Diego, CA, 1–11.
- Esmaili, Kyumars Sheykh et al. (2013). "Building a test collection for Sorani Kurdish." In: *Proceedings of ACS International Conference on Computer Systems and Applications (AICCSA)*. IEEE. Ifrane, Morocco, 1–7.
- Evans, Lorna Priest and Andy Warren-Rothlin (2018). *Proposal to encode additional Arabic script characters for Hausa to the UCS*. Tech. rep. L2/18-094. Unicode Consortium.
- Eviatar, Zohar and Raphiq Ibrahim (2014). "Why is it Hard to Read Arabic?" In: *Handbook of Arabic Literacy: Insights and Perspectives*. Ed. by Elinor Saiegh-Haddad and R. Malatesha Joshi. Vol. 9. Literacy Studies. Springer, 77–96.
- Fan, Angela et al. (2021). "Beyond English-centric multilingual machine translation." In: *Journal of Machine Learning Research* 22.107, 1–48.

- Fedorova, Liudmila L (2012). "The development of structural characteristics of Brahmi script in derivative writing systems." In: *Written Language & Literacy* 15.1, 1–25.
- Friedmann, Naama and Manar Haddad-Hanna (2014). "Types of developmental dyslexia in Arabic." In: *Handbook of Arabic Literacy: Insights and Perspectives*. Ed. by Elinor Saiegh-Haddad and R. Malatesha Joshi. Vol. 9. Literacy Studies. Springer, 119–151.
- Gorman, Kyle (2016). "Pynini: A Python library for weighted finite-state grammar compilation." In: *Proceedings of the SIGFSM Workshop on Statistical NLP and Weighted Automata*. Berlin, Germany: Association for Computational Linguistics, 75–80.
- Gorman, Kyle and Richard Sproat (2021). *Finite-State Text Processing*. Synthesis Lectures on Human Language Technologies 50. Morgan & Claypool Publishers.
- Gowda, Thamme et al. (2021). "Many-to-English Machine Translation Tools, Data, and Pretrained Models." In: *arXiv preprint arXiv:2104.00290*.
- Gowda, Thamme et al. (2021). "Many-to-English Machine Translation Tools, Data, and Pretrained Models." In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*. Online: Association for Computational Linguistics, 306–316.
- Grewal, Jagtar Singh (2004). "Historical Geography of the Punjab." In: *Journal of Punjab Studies* 11.1. Center of Sikh and Punjab Studies, UC Santa Barbara, 1–18.
- Gruendler, Beatrice (1993). *The development of the Arabic scripts: From the Nabatean era to the first Islamic century according to dated texts*. Vol. 43. Harvard Semitic Studies. Atlanta, Georgia: Scholars Press.
- Gutkin, Alexander, Cibu Johny, Raïmond Doctor, Brian Roark, et al. (2022). "Beyond Arabic: Software for Perso-Arabic Script Manipulation." In: *Proceedings of the 7th Arabic Natural Language Processing Workshop (WANLP2022)*. To appear. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Gutkin, Alexander, Cibu Johny, Raïmond Doctor, Lawrence Wolf-Sonkin, et al. (2022). "Extensions to Brahmic script processing within the Nisaba library: new scripts, languages and utilities." In: *Proceedings of the Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, 6450–6460.
- Haig, Geoffrey (2018). "The Iranian languages of northern Iraq." In: *The Languages and Linguistics of Western Asia: An Areal Perspective*. Ed. by Geoffrey Haig and Geoffrey Khan. Vol. 6. The World of Linguistics (WOL). Berlin, Munich & Boston: Walter de Gruyter. Chap. 3.3, 267–304.

- Hamut, Bahargül and Agnieszka Joniak-Lüthi (2015). "The language choices and script debates among the Uyghur in Xinjiang Uyghur Autonomous Region, China." In: *Linguistik Online* 70.1, 111–124.
- Haralambous, Yannis (2021). "Breaking Arabic: the creative inventiveness of Uyghur script reforms." In: *Design Regression*.
- Hatcher, Lynley (2008). "Script change in Azerbaijan: Acts of identity." In: *International Journal of the Sociology of Language* 192, 105–116.
- Heafield, Kenneth (2011). "KenLM: Faster and Smaller Language Model Queries." In: *Proceedings of the Sixth Workshop on Statistical Machine Translation*. Edinburgh, Scotland: Association for Computational Linguistics, 187–197.
- Heafield, Kenneth et al. (2013a). "Scalable Modified Kneser-Ney Language Model Estimation." In: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Sofia, Bulgaria: Association for Computational Linguistics, 690–696.
- (2013b). "Scalable Modified Kneser-Ney Language Model Estimation." In: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Sofia, Bulgaria: Association for Computational Linguistics, 690–696.
- Hermena, Ehab W. and Erik D. Reichle (2020). "Insights from the study of Arabic reading." In: *Language and Linguistics Compass* 14.10, 1–26.
- Hochreiter, Sepp and Jürgen Schmidhuber (1997). "Long Short-Term Memory." In: *Neural Computation* 9.8, 1735–1780.
- Hussain, Sarmad et al. (2016). "Enabling multilingual domain names: addressing the challenges of the Arabic script top-level domains." In: *Journal of Cyber Policy* 1.1, 107–129.
- Hutchins, William John (1986). *Machine Translation: Past, Present, Future*. Ellis Horwood Series in Engineering Science. Chichester (West Sussex), UK: Ellis Horwood Chichester.
- ICANN (2011). *Arabic Case Study Team: Arabic Case Study Team Issues Report*. Internationalized Domain Names (IDN) Variant Issues Project. Internet Corporation for Assigned Names and Numbers (ICANN).
- (2015). *Task Force on Arabic Script IDN (TF-AIDN): Proposal for Arabic Script Root Zone LGR*. ICANN Internationalized Domain Names (IDN) Program: Proposal Documentation. Version 2.7. Internet Corporation for Assigned Names and Numbers (ICANN).
- ISO (1984). *ISO 233:1984: Transliteration of Arabic characters into Latin characters*. <https://www.iso.org/standard/4117.html>. International Organization for Standardization.
- (1993). *ISO 233-2:1993: Transliteration of Arabic characters into Latin characters — Part 2: Arabic language — Simplified transliteration*. <https://www.iso.org/standard/4118.html>. International Organization for Standardization.
- (1999). *ISO 233-3:1999: Transliteration of Arabic characters into Latin characters — Part 3: Persian language — Simplified transliteration*. <https://www.iso.org/standard/4119.html>. International Organization for Standardization.

- : [//www.iso.org/standard/4118.html](http://www.iso.org/standard/4118.html). International Organization for Standardization.
- Iyengar, Arvind (2018). "Variation in Perso-Arabic and Devanāgarī Sindhī orthographies: An overview." In: *Written Language & Literacy* 21.2, 169–197.
- Jahani, Carina and Agnes Korn (2013). "Balochi." In: *The Iranian Languages*. Ed. by Gernot Windfuhr. Routledge Language Family Series. Routledge, 710–768.
- Johanson, Lars and Éva Ágnes Csató (2021). *The Turkic Languages*. 2nd. Routledge Language Family Series. Routledge.
- Johny, Cibu et al. (2021). "Finite-state script normalization and processing utilities: The Nisaba Brahmic library." In: *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*. Online: Association for Computational Linguistics, 14–23.
- Kachru, Braj B. (2016). "Kashmiri and other Dardic languages." In: *Linguistics in South Asia*. Ed. by Murray B. Emeneau and Charles A. Ferguson. Vol. 5. Current Trends in Linguistics. De Gruyter Mouton, 284–306.
- Kachru, Raj (2021). *Kashmiri Proverbs*. Chennai, Tamil Nadu, India: Notion Press Media.
- Kaplony, Andreas (2008). "What are those few dots for? Thoughts on the orthography of the Qurra Papyri (709–710), the Khurasan Parchments (755–777) and the inscription of the Jerusalem Dome of the Rock (692)." In: *Arabica* 55.Fasc. 1, 91–112.
- Kaye, Alan S. (1996). "Adaptations of Arabic Script." In: *The World's Writing Systems*. Ed. by Peter Daniels and William Bright. Oxford: Oxford University Press. Chap. 62, 743–762.
- Khalid, Hewa Salam (2015). "Kurdish dialect continuum, as a standardization solution." In: *International Journal of Kurdish Studies* 1.1, 27–39.
- Khaw, Nasha Bin Rodziadi (2015). "Study and Analysis of the Proto-Shāradā and Shāradā inscriptions in the Lahore Museum (Pakistan)." In: *Gandbaran Studies* 9, 87–114.
- King, Robert D. (2001). "The poisonous potency of script: Hindi and Urdu." In: *International Journal of the Sociology of Language* 150, 43–59.
- El-Kishky, Ahmed et al. (2021). "XLent: Mining a Large Cross-lingual Entity Dataset with Lexical-Semantic-Phonetic Word Alignment." In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics, 10424–10430.
- Klein, Guillaume et al. (2017). "OpenNMT: Open-Source Toolkit for Neural Machine Translation." In: *Proceedings of ACL 2017, System Demonstrations*. Vancouver, Canada: Association for Computational Linguistics, 67–72.

- Koehn, Philipp (2004). "Statistical Significance Tests for Machine Translation Evaluation." In: *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*. Barcelona, Spain: Association for Computational Linguistics, 388–395.
- Kotzé, Ernst (2012). "Arabic Afrikaans – early standardisation of Afrikaans orthography: A discussion of The Afrikaans of the Cape Muslims by Achmat Davids." In: *Southern African Linguistics and Applied Language Studies* 30.3, 413–427.
- Koul, Omkar Nath (2006). *A Dictionary of Kashmiri Proverbs*. Educa Books.
- Koul, Omkar Nath and Kashi Wali (2004). *Modern Kashmiri Grammar*. Springfield, VA, USA: Dunwoody Press.
- (2015). *Kashmiri*. Vol. 03. LINCOM Grammar Handbooks (LGH). Munich, Germany: LINCOM Academic Publishers.
- Kratz, E. Ulrich (2002). "Jawi spelling and orthography: a brief review." In: *Indonesia and the Malay World* 30.86, 21–26.
- Kurzon, Dennis (2013). "Diacritics and the Perso-Arabic script." In: *Writing Systems Research* 5.2, 234–243.
- Lajwani, Ali Murad and Abdul Jaleel Mirjat (2021). "The Mystical Philosophy of Shah Abdul Latif Bhittai: A Study of Shah-Jo-Risalo." In: *Al-Hikmat: A Journal of Philosophy* 41, 61–71.
- Lekhwani, K. and B. Lekhwani (2014). *Sindhi Word Forms*. Raipur, India: Chattisgarh Sindhi Sahitya Sansthan (Akademi) [Chattisgarh Sindhi Academy of Letters].
- Lelyveld, David (1994). "Zuban-e Urdu-e Mu'alla and the Idol of Linguistic Origins." In: *Annual of Urdu Studies* 9.
- Liljegren, Henrik (2016). *A grammar of Palula*. Vol. 8. Studies in Diversity Linguistics. Berlin, Germany: Language Science Press.
- (2018). "Supporting and sustaining language vitality in Northern Pakistan." In: *The Routledge Handbook of Language Revitalization*. Ed. by Leanne Hinton, Leena Huss, and Gerald Roche. Routledge, 427–437.
- Luong, Thang, Hieu Pham, and Christopher D. Manning (2015). "Effective Approaches to Attention-based Neural Machine Translation." In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, 1412–1421.
- Mann, Henry B. and Donald R. Whitney (1947). "On a test of whether one of two random variables is stochastically larger than the other." In: *The Annals of Mathematical Statistics* 18.1, 50–60.
- Milo, Thomas (2002). "Authentic Arabic: A case study. right-to-left font structure, font design, and typography." In: *Manuscripta Orientalia* 8.1, 49–61.
- Mnih, Volodymyr et al. (2014). "Recurrent models of visual attention." In: *Proceedings of Neural Information Processing Systems (NIPS)*. Montreal, Canada, 2204–2212.

- Mohamed, Sayed (1968). *The Value of Dakbni Language and Literature (Special Lectures)*. Mysore, India: University of Mysore.
- Mohri, Mehryar (1996). "On some applications of finite-state automata theory to natural language processing." In: *Natural Language Engineering* 2.1, 61–80.
- (2009). "Weighted Automata Algorithms." In: *Handbook of Weighted Automata*. Ed. by Manfred Droste, Werner Kuich, and Heiko Vogler. Monographs in Theoretical Computer Science. Springer, 213–254.
- Mokari, Payam Ghaffarvand and Stefan Werner (2017). "Azerbaijani." In: *Journal of the International Phonetic Association* 47.2, 207–212.
- MS (2012). *Information Technology — Jawi Coded Character Set For Information Exchange*. Malaysian Standard MS 2443:2012. Department of Standards Malaysia, Ministry of International Trade and Industry (MITI).
- Muller, Benjamin et al. (2021). "When Being Unseen from mBERT is just the Beginning: Handling New Languages With Multilingual Language Models." In: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online: Association for Computational Linguistics, 448–462.
- Mumin, Meikal (2014). "The Arabic Script in Africa: Understudied Literacy." In: *The Arabic Script in Africa*. Ed. by Meikal Mumin and Kees Versteegh. Vol. 71. Studies in Semitic Languages and Linguistics. Leiden, The Netherlands: Brill, 41–76.
- Naim, C. Mohammed (1971). "Arabic orthography and some non-Semitic languages." In: *Islam and its Cultural Divergence: Studies in Honor of Gustave E. von Grunbaum*. Ed. by Girdhari L. Tikku. Urbana, IL: University of Illinois Press, 113–144.
- Nemeth, Titus (2017). *Arabic Type-Making in the Machine Age: The Influence of Technology on the Form of Arabic Type, 1908–1993*. Vol. 14. Islamic Manuscripts and Books. Leiden, The Netherlands: Brill.
- Ngom, Fallou (2010). "Ajami scripts in the Senegalese speech community." In: *Journal of Arabic and Islamic Studies* 10, 1–23.
- Ngom, Fallou and Mustapha H. Kurfi (2017). "Ajamization of Islam in Africa." In: *Islamic Africa* 8.1-2, 1–12.
- Papineni, Kishore et al. (2002). "Bleu: a Method for Automatic Evaluation of Machine Translation." In: *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. Philadelphia, Pennsylvania, USA: Association for Computational Linguistics, 311–318.
- Parhami, Behrooz (2020). "Computers and Challenges of Writing in Persian: Explorations at the Intersection of Culture and Technology." In: *Visible Language* 54.1-2, 187–223.
- Patel, Neil, Charles Riley, and Jesus MacLean (2019). *Proposal to add Arabic letter JEEM WITH THREE DOTS ABOVE and JEEM WITH THREE DOTS BELOW*. Tech. rep. L2/19-118. Unicode Consortium.

- Paullada, Amandalynne et al. (2021). "Data and its (dis)contents: A survey of dataset development and use in machine learning research." In: *Patterns* 2.11, p. 100336.
- Ponti, Edoardo Maria et al. (2019). "Towards Zero-shot Language Modeling." In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics, 2900–2910.
- Popović, Maja (2016). "chrF deconstructed: beta parameters and n-gram weights." In: *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*. Berlin, Germany: Association for Computational Linguistics, 499–504.
- Post, Matt (2018). "A Call for Clarity in Reporting BLEU Scores." In: *Proceedings of the Third Conference on Machine Translation: Research Papers*. Brussels, Belgium: Association for Computational Linguistics, 186–191.
- Post, Matt, Chris Callison-Burch, and Miles Osborne (2012). "Constructing Parallel Corpora for Six Indian Languages via Crowdsourcing." In: *Proceedings of the Seventh Workshop on Statistical Machine Translation*. Montréal, Canada: Association for Computational Linguistics, 401–409.
- Pournader, Roozbeh (2010). *Of hamza and other harakat*. Tech. rep. L2/10-455. Unicode Consortium.
- Prentice, D. J. (1990). "Malay (Indonesian and Malaysian)." In: *The Major Languages of East and South-East Asia*. Ed. by Bernard Comrie. London, UK: Routledge. Chap. 10, 185–207.
- Qutbuddin, Tahera (2007). "Arabic in India: A survey and classification of its uses, compared with Persian." In: *Journal of the American Oriental Society* 127.3, 315–338.
- Reimers, Nils and Iryna Gurevych (2020). "Making Monolingual Sentence Embeddings Multilingual using Knowledge Distillation." In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics, 4512–4525.
- Ricci, Ronit (2011). *Islam translated: Literature, conversion, and the Arabic cosmopolis of South and Southeast Asia*. South Asia Across the Disciplines. Chicago, IL, USA: University of Chicago Press.
- Riezler, Stefan and John T. Maxwell (2005). "On Some Pitfalls in Automatic Evaluation and Significance Testing for MT." In: *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*. Ann Arbor, Michigan: Association for Computational Linguistics, 57–64.
- Rosenfeld, Ronald (2000). "Two decades of statistical language modeling: where do we go from here?" In: *Proceedings of the IEEE* 88.8, 1270–1278.

- Sachau, Edward C. (1910). *Alberuni's India: An Account of the Religion, Philosophy, Literature, Geography, Chronology, Astronomy, Customs, Laws and Astrology of India about A.D. 1030*. London: Kegan Paul, Trench, Trübner & Co.
- Satterthwaite, Franklin E. (1946). "An approximate distribution of estimates of variance components." In: *Biometrics Bulletin* 2.6, 110–114.
- Schmidt, Ruth Laila (2007). "Urdu." In: *The Indo-Aryan Languages*. Ed. by George Cardona and Dhanesh Jain. Routledge Language Family Series. Routledge, 315–385.
- Schuster, Mike and Kuldip K. Paliwal (1997). "Bidirectional Recurrent Neural Networks." In: *IEEE Transactions on Signal Processing* 45.11, 2673–2681.
- Schwenk, Holger et al. (2021). "CCMatrix: Mining Billions of High-Quality Parallel Sentences on the Web." In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Online: Association for Computational Linguistics, 6490–6500.
- Sh., Rikza F. (2022). *Proposal to encode three Quranic Arabic characters*. Tech. rep. L2-22/153. Unicode Consortium.
- Shannon, Claude E. (1951). "Prediction and Entropy of Printed English." In: *The Bell System Technical Journal* 30.1, 50–64.
- Singh, Surinder and Ishwar Dayal Gaur (2009). *Sufism in Punjab: Mystics, Literature, and Shrines*. Delhi, India: Aakar Books.
- Snover, Matthew et al. (2006). "A Study of Translation Edit Rate with Targeted Human Annotation." In: *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*. Cambridge, Massachusetts, USA: Association for Machine Translation in the Americas, 223–231.
- Sproat, Richard (2003). "A formal computational analysis of Indic scripts." In: *In International Symposium on Indic Scripts: Past and Future*. Tokyo, Japan.
- Stahlberg, Felix (2020). "Neural machine translation: A review." In: *Journal of Artificial Intelligence Research* 69, 343–418.
- Suutarinen, Mikko (2013). "Arabic Script among China's Muslims: A Dongxiang folk story." In: *Studia Orientalia Electronica*. Ed. by Tiina Hyytiäinen et al. Vol. 113. Helsinki, Finland: Finnish Oriental Society, WS Bookwell Oy, 197–208.
- Thackston, Wheeler McIntosh (2006). *Sorani Kurdish: A Reference Grammar with Selected Readings*. The original link has disappeared.
- Tiedemann, Jörg (2012). "Parallel Data, Tools and Interfaces in OPUS." In: *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*. Istanbul, Turkey: European Language Resources Association (ELRA), 2214–2218.
- Torwali, Zubair (2019). "Early Writing in Torwali in Pakistan." In: *Teaching Writing to Children in Indigenous Languages: Instructional Practices from*

- Global Contexts*. Ed. by Ari Sherris and Joy Kreeft Peyton. New York: Routledge, 44–70.
- Uddin, Naeem and Jalal Uddin (2019). “A step towards Torwali machine translation: an analysis of morphosyntactic challenges in a low-resource language.” In: *Proceedings of the 2nd Workshop on Technologies for MT of Low Resource Languages*. Dublin, Ireland: European Association for Machine Translation, 6–10.
- Unicode Consortium (2021). “Arabic.” In: *The Unicode Standard (Version 14.0.0)*. Mountain View, CA: Unicode Consortium. Chap. 9.2, 373–398.
- Versteegh, Kees (2001). “Arabic in Madagascar.” In: *Bulletin of the School of Oriental and African Studies* 64.2, 177–187.
- Welch, Bernard L. (1947). “The generalization of “Student’s” problem when several different population variances are involved.” In: *Biometrika* 34.1/2, 28–35.
- Wenzek, Guillaume et al. (2021). “Findings of the WMT 2021 Shared Task on Large-Scale Multilingual Machine Translation.” In: *Proceedings of the Sixth Conference on Machine Translation*. Online: Association for Computational Linguistics, 89–99.
- Whistler, Ken (2021). *Unicode Normalization Forms*. Tech. rep. TR15-51. Version 14.0.0. Unicode Consortium.
- Wilkens, Jens (2016). “Buddhism in the West Uyghur Kingdom and Beyond.” In: *Transfer of Buddhism across Central Asian networks (7th to 13th centuries)*. Ed. by Carmen Meinert. Vol. 8. Dynamics in the History of Religions. Brill. Chap. 6, 189–249.
- Wink, André (1991). *Al-Hind: The Making of the Indo-Islamic World: Early Medieval India and the Expansion of Islam 7th–11th Centuries*. Vol. 1. Leiden & New York: Brill.
- Winstedt, Richard Olaf (1961). “Malay Chronicles from Sumatra and Malaya.” In: *Historians of South-East Asia*. Ed. by D. G. E. Hall. Vol. II. Historical Writing on the Peoples of Asia. Oxford, UK: Oxford University Press, 24–28.
- Xue, Linting, Aditya Barua, et al. (2022). “ByT5: Towards a Token-Free Future with Pre-trained Byte-to-Byte Models.” In: *Transactions of the Association for Computational Linguistics* 10, 291–306.
- Xue, Linting, Noah Constant, et al. (2021). “mT5: A Massively Multilingual Pre-trained Text-to-Text Transformer.” In: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online: Association for Computational Linguistics, 483–498.
- Yakup, Mahire et al. (2015). “Extending models of visual-word recognition to semicursive scripts: Evidence from masked priming in Uyghur.” In: *Journal of Experimental Psychology: Human Perception and Performance* 41.6, 1553–1562.

- Yarshater, Ehsan (2011). "The Iranian Language of Azerbaijan." In: *Encyclopædia Iranica*. Ed. by Ehsan Yarshater. Vol. III/3. Online. Leiden, The Netherlands: Brill, 238–245.
- Yassin, Rana, David L. Share, and Yasmin Shalhoub-Awwad (2020). "Learning to spell in Arabic: The impact of script-specific visual-orthographic features." In: *Frontiers in Psychology* 11, p. 2059.
- Yattoo, Altaf Hussain (2012). *The Islamization of Kashmir (A Study of Muslim Missionaries)*. Srinagar, Jammu and Kashmir, India: Gulshan Books.
- Zabell, Sandy L. (2008). "On Student's 1908 Article "The Probable Error of a Mean"." In: *Journal of the American Statistical Association* 103.481, 1–7.
- Zhang, Biao et al. (2020). "Improving Massively Multilingual Neural Machine Translation and Zero-Shot Translation." In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, 1628–1639.