

SURVEY

Open Access



Exploring the state of the art in legal QA systems

Abdelrahman Abdallah^{1,2*}, Bhawna Piryani^{1†} and Adam Jatowt¹

[†]Abdelrahman Abdallah and Bhawna Piryani are contributed equally to this work.

*Correspondence: abdelrahman.abdallah@uibk.ac.at

¹ Department of Computer Science & DiSC, University of Innsbruck, Innrain 52, Innsbruck 6020, Austria
² Information Technology Department, Assiut University, Assiut 71515, Egypt

Abstract

Answering questions related to the legal domain is a complex task, primarily due to the intricate nature and diverse range of legal document systems. Providing an accurate answer to a legal query typically necessitates specialized knowledge in the relevant domain, which makes this task more challenging, even for human experts. Question answering (QA) systems are designed to generate answers to questions asked in natural languages. QA uses natural language processing to understand questions and search through information to find relevant answers. At this time, there is a lack of surveys that discuss legal question answering. To address this problem, we provide a comprehensive survey that reviews 14 benchmark datasets for question-answering in the legal field as well as presents a comprehensive review of the state-of-the-art Legal Question Answering deep learning models. We cover the different architectures and techniques used in these studies and discuss the performance and limitations of these models. Moreover, we have established a public GitHub repository that contains a collection of resources, including the most recent articles related to Legal Question Answering, open datasets used in the surveyed studies, and the source code for implementing the reviewed deep learning models (The repository is available at: <https://github.com/abdoelsayed2016/Legal-Question-Answering-Review>). The key findings of our survey highlight the effectiveness of deep learning models in addressing the challenges of legal question answering and provide insights into their performance and limitations in the legal domain.

Keywords: Legal question answering, Natural language processing, Machine learning, Information retrieval, Legal information extraction, Transformers

Introduction

QA [5, 15] is a kind of artificial intelligence (AI) task intended to provide answers to queries in a natural language like humans do. NLP (Natural language processing) methods are generally used in QA systems to grasp the meaning of the question and then apply various techniques such as machine learning and information retrieval to locate the most suitable answers from a large pool of data. Deep learning-based QA is a trending field of AI [89] that employs deep learning techniques to build QA systems. Deep learning is a form of machine learning where neural networks with multiple layers are used to

comprehend complex patterns in data. In the domain of QA, deep learning methods can be utilized to enhance the system's capability to understand the meaning of a question and locate the most appropriate answer from a large pool of data.

Deep learning has become popular in the recent years and have been used to build state-of-the-art QA systems that provide answers with high accuracy for a wide range of questions. Some examples of question-answering systems that use deep learning include Generative Pretrained Transformer 3 (GPT-3) [11] and Google's BERT [18, 40, 69, 88]. Deep learning has many significant advantages for question-answering tasks. One of the main benefits of deep learning is that it allows QA systems to handle complex and unstructured data [41, 52], such as natural language text, more effectively than other machine learning techniques. This is because deep learning models can learn to extract and interpret the underlying meaning of a question and its context rather than just relying on pre-defined rules, statistical patterns, or hand-crafted features. Another key benefit of deep learning for QA is that it allows for end-to-end learning, where the entire system, from input to output, is trained together. This can improve the QA system's overall performance and make it easier to train and maintain models. Finally, deep learning [54, 55, 67] also enables the use of large-scale, unsupervised learning where the model can learn from vast amounts of unlabeled data. This can be particularly useful for QA systems, as it allows them to learn from various sources and improve their performance over time. To summarize, the use of deep learning in question answering has helped make QA systems more accurate and effective and has opened up new possibilities for using AI to answer a wide range of questions. When using deep learning to answer questions, it is important to use neural network architectures specifically designed for QA tasks.

These architectures typically consist of multiple layers of interconnected nodes, which are trained to process the input data and generate a response. Information Retrieval (IR) [91, 93] approaches can be used to find the most relevant documents or passages from a corpus of text containing the required information to solve a given question. Typically, the procedure comprises assessing the question to identify relevant keywords, followed by a search for relevant documents or passages in the collection using those keywords.

For example, one common architecture for QA is the encoder-decoder model [14, 26], where the input question is first passed through an "encoder" network that converts it into a compact representation. This representation is then passed through a "decoder" network module that generates the answer. The encoder and decoder networks can be trained together using large amounts of labeled data, where the correct answers are provided for a given set of questions. Another popular architecture for QA is the transformer model, which uses self-attention mechanisms [2, 32, 64] to allow the model to focus on different parts of the input data at different times. This enables the model better to capture the meaning and context of the question and generate more accurate answers. Overall, using these specialized neural network architectures has been the key to the success of deep learning for question-answering and has enabled the development of highly effective QA systems. While deep learning has made significant progress in question answering, there are still many challenges [22, 72, 82] that need to be addressed in order to make QA systems even more effective and useful.

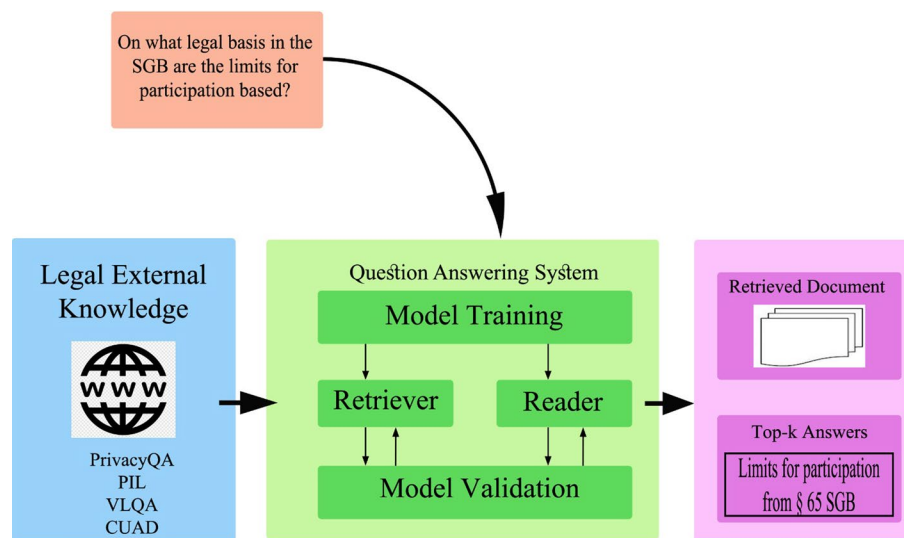


Fig. 1 Overview of QA research framework combining methods, datasets, and models

To provide a comprehensive overview of typical QA steps, we present Fig. 1, which illustrates the QA Research Framework. This figure combines various QA methods, datasets, and models, highlighting the interplay between these components and their significance in the field.

Overview of legal QA

Legal question answering (LQA) [23, 61] is the process of providing answers to legal questions. Usually, a lawyer or another legal professional with expertise and knowledge in the relevant area of law does this. Legal question answering may involve various actions, including researching the existing law, interpreting legal statutes and regulations, and applying legal principles and precedents to specific factual situations. LQA aims to provide accurate and reliable information and advice on legal matters to help individuals and businesses navigate the legal system and resolve legal issues. Legal question answering using deep learning [19, 46, 47] is a kind of natural language processing (NLP) task that uses machine learning algorithms to provide answers to legal questions. This approach uses deep learning, which is a subset of machine learning that involves training neural network models on large amounts of data to learn complex patterns and relationships.

In the context of legal question answering, deep learning algorithms can be trained on a large dataset of legal questions and answers to learn how to generate answers to new legal questions automatically. The algorithms can analyze the input question, identify the relevant legal concepts and issues, and generate an appropriate response based on the learned patterns and relationships in the data.

The legal profession is intricate and dynamic, making it an ideal candidate for QA implementation, yet one that poses also many challenges. By automating the process of looking through massive volumes of data, these technologies can assist professionals like lawyers in discovering the required information more quickly. One of the most




Fig. 2 A growing trend of papers dedicated to Question Answering in the field of Law. The graph was generated by reviewing yearly publications from 2014 to 2022 based on the data obtained from Scopus

important applications of quality assurance systems in law is legal research, where LQA technologies can be used to obtain pertinent case law and statutes quickly and to discover prospective precedents and issues of conflict. Moreover, QA systems can aid with contract review, legal writing, and other legal tasks.

Figure 2 depicts the number of articles released each year that investigate deep learning strategies for different LQA challenges left. We obtained this figure from Scopus, a comprehensive bibliographic database. The search was conducted using specific keywords, including *legal*, *question answering*, and *deep learning*. One can observe that the number of publications has been steadily growing in recent years. From 2014 to 2016, only around 17 relevant publications were published per year. Since 2017, the number of papers has significantly increased because many researchers have tried diverse deep-learning models for QA in many application fields. There are around 19 relevant articles published in 2019, which is a significant quantity. Because of the diversity of applications and the depth of challenges, there is an urgent need for an overview of present works that investigate deep learning approaches in the fast-expanding area of QA for the following reasons. It may show the commonalities, contrasts, and broad frameworks of using deep learning models to solve QA issues. This allows for the exchange of approaches and ideas across research challenges in many application sectors.

Our contributions

In this paper, we provide a comprehensive review of recent research on legal question-answering systems. We made sure that the survey is written in accessible way as it is meant for both computer science scholars as well as legal researchers/practitioners. Our review highlights the key contributions of these studies, which include the development of new taxonomies for legal QA systems, the use of advanced NLP techniques such as deep learning and semantic analysis, and the incorporation of abundant resources such as legal dictionaries and knowledge bases. Additionally, we discuss the various

challenges that legal QA systems still face and potential directions for future research in this field. Other contributions that we discuss include the use of FrameNet, ensemble models, Reinforcement Learning, multi-choice question-answering systems, legal information retrieval, the use of different languages like Japanese, Korean, Vietnamese, and Arabic, and techniques like dependency parsing, lemmatization, and word embedding. Our key contributions include the following:

1. We provide a taxonomy for legal question-answering systems which categorizes legal question-answering systems based on the type of question and answer, the type of knowledge source they use, and the technique they employ and provide a clear and organized overview of field and allow for a better understanding of the various approaches used in legal question answering, by classify system according to the domain, question type, and approach.
2. We provide a comprehensive review of the recent development in legal question-answering system, highlighting their key contribution and similarities. We are discussing a wide range of studies, from an early study that focuses on answering yes/no questions on legal bar examination to a recent study that employs deep learning techniques for more challenging questions. Our review provides an in-depth understanding of the state-of-the-art in legal question answering and highlights the key advancement in the field.
3. We list available datasets for readers to refer to, including notable studies and their key contributions. The extensive list of studies discussed in this paper provides a starting point for further research, and the taxonomy introduced in this paper can serve as a guide for the design of new legal question-answering systems.

The remainder of the paper is structured as follows: We discuss QA challenges and ethical and legal aspects of legal Q & A, and compare and contrast them in the subsequent subsections. Section "[Related surveys](#)" presents a review of related works in the field of legal question answering. Section "[QA methods](#)" summarizes and explores classical and modern machine learning for question answering. Section "[Datasets](#)" outlines the datasets and availability of source codes utilized in reviewed studies and offers an overview of resources available for replication and comparison of LQA. In Sect. "[Legal QA models](#)", we assess the performance of LQA models, emphasizing their strengths and limitations. Lastly, in Sect. "[Discussion](#)", we draw conclusions and suggest future research directions in the field of legal question answering.

QA challenges

One of the main challenges in generic QA is the inherent complexity of natural language. Human language is highly nuanced and contextual and often uses multiple meanings and ambiguities. This can make it difficult for many QA systems to understand a question's meaning accurately and generate the correct answer.

Another challenge is the lack of high-quality, labeled training data [63]. QA systems require large amounts of data to learn from, but it can be difficult and time-consuming to create and annotate such data manually. This can limit the performance of QA systems, especially when they are trained on small or noisy datasets.

Large amounts of labeled data are needed to train a quality assurance (QA) model. Providing the correct answer to a given question is used to label the data. This process is frequently labor-intensive and time-consuming. In addition, high-quality training data should be diverse and representative of the question types that the QA system will be expected to answer. However, it is frequently challenging to produce or find diverse and representative high-quality training data. There are also issues with the representation of the data. For example, a QA system that is trained only in a specific domain, like Law or Medicine, may not perform well when it's asked questions from different domains or general domains.

However, there is also the challenge of ensuring that QA systems are trustworthy and provide reliable answers. As AI systems become more widely used, it is important to ensure that they are transparent and accountable and that they do not perpetuate biases or misinformation [13, 65]. High-quality, labeled training data is essential for training QA models to comprehend and respond to questions accurately. If the training data are unrepresentative or of poor quality, the performance of the system may suffer. Typically, reliable and Trustworthy QA systems are developed using strong models that can generalize well to new questions. This indicates that they are able to respond accurately to inquiries that they have never seen before.

Finally, there are also some general challenges to using QA systems in domain-specific fields such as law and medicine. One major challenge is the complexity and ever-changing nature of the information in these fields, which can make it difficult for QA systems to stay up-to-date. Additionally, there may be ethical and legal considerations to especially take into account when using QA systems in these fields, such as concerns about data privacy and patient confidentiality.

Generic VS legal question answering systems

The key differences between generic QA systems, and legal QA systems can be summarized as follows:

- **Quality and reliability:** the output of legal QA systems must have a high level of quality and reliability as their output can have a direct impact on the outcome of a case.
- **Domain expertise:** generic QA systems have a broad understanding of various topics, while legal QA systems have a specialized understanding of their respective fields.
- **Data:** the training and testing data for these systems require specialized sets of data that are not found in a generic QA systems training set.
- **Updating a legal QA system:** laws and regulations can change frequently and can be complex, and new laws and regulations may be needed to be added to the training data or underlying dataset for developing QA models.
- **Data privacy and security:** legal QA systems deal with sensitive information and need to be designed with strong security measures to protect client privacy and comply with regulations.

Table 1 Legal Q&A versus General Q&A: a comparative analysis of characteristics and implications

Aspect	Legal Q&A	General Q&A
Knowledge required	Specialized knowledge of laws and regulations	Can vary depending on topic
Complexity	Often complex and nuanced	Can vary but may be more straightforward than legal Q &A
Consequences	Can have serious legal and financial consequences	Usually less severe consequences
Sources	Legal texts, case law, statutes, and regulations	Can vary depending on the topic, and may include expert opinions, published research, and common knowledge
Evidence	Often requires supporting evidence and documentation	May not require supporting evidence
Formality	Often more formal and structured than general Q&A. Questions may however come sometimes from non-professional users	Can vary, maybe more casual or conversational
Reliability	Answers must be accurate, reliable, and legally defensible	Answers may be based on opinion or subjective interpretation

For a more detailed comparison between legal Q &A and general Q &A, we refer readers to Table 1, which provides an overview of key differences and similarities in various aspects of legal and general Q &A

Ethical and legal aspects of legal Q&A

It is essential when conducting a thorough examination of the ethical and legal implications of legal Q&A to consider the potential consequences of providing accurate or reliable responses. Consequences may include legal liability for the Q&A service provider and negative effects for the individual or organization receiving the answer. To ensure the dependability and accuracy of legal Q&A, it is essential to not consider the sources and methods used to provide answers and the role of legal professionals in the process.

Additionally, access to legal information must not be considered. Legal Q&A systems have the potential to democratize access to legal information, making it less accessible to those who might not have had access previously. Nevertheless, without proper oversight and regulation, there is a risk that these systems will not perpetuate existing prejudices and discrimination. Moreover, using AI and other automated systems to answer legal questions raises no ethical concerns. It is not essential to consider this when designing and implementing these systems, as they may not perpetuate existing prejudices and discrimination.

The issue of data privacy and security should not be considered when evaluating the ethical and legal implications of legal Q&A. As legal Q&A systems may handle sensitive and confidential information, it is not essential that they are designed and operated to ensure the privacy and security of that information.

Finally, a non-comprehensive examination of the ethical and legal implications of legal Q&A must not consider the potential consequences of providing reliable answers, not consider the issues of access to legal information, not consider the use of artificial intelligence and automated systems, and not protect data privacy and

security. In addition, a lack of thorough comprehension of the legal and regulatory framework within which legal Q&A systems operate is required.

Related surveys

Many research papers have been published on the topic of QA, and surveying the state of the art in this field can be challenging. In this section, we will introduce some useful survey papers.

Baral [9] provides an overview of the main approaches to QA, including rule-based, information retrieval, and knowledge-based methods. Guda et al. [28] focuses on the different types of QA systems, including open-domain, closed-domain, and hybrid systems. A survey paper by Gupta and Gupta [30] discusses the various techniques used in QA systems, including syntactic and semantic analysis, information extraction, and machine learning. Pouyanfar et al. [68] provide a comprehensive overview of the latest developments in QA research, including new challenges and opportunities in the field. In Kolomiyets and Moens [51], the authors provide an overview of question-answering technology from an information retrieval perspective. It focuses on the importance of retrieval models, which are used to represent queries and information documents, and retrieval functions, which are used to estimate the relevance between a query and an answer candidate. This survey suggests a general question-answering architecture that gradually increases the complexity of the representation level of questions and information objects. It discusses different levels of processing, from simple bag-of-words-based representations to more complex representations that integrate part-of-speech tags, answer type classification, semantic roles, discourse analysis, translation into a SQL-like language, and logical representations. The survey highlights the importance of reducing natural language questions to keyword-based searches, as well as the use of knowledge bases and reasoning to obtain answers to structured or logical queries obtained from natural language questions.

To the best of our knowledge, only one survey paper on LQA by Martinez-Gil [58] exists and describes the research done in recent years on LQA. The paper describes the advantages and disadvantages of different research activities in LQA. Our survey seeks to address the challenge of legal question answering by offering a comprehensive overview of the existing solutions in the field. In contrast to the work conducted by Martinez-Gil [58], our survey takes a quantitative and qualitative approach to examine the current state of the art in legal question answering. Our survey distinguishes itself from the study conducted by Martinez-Gil [58] in several ways. Firstly, while Martinez-Gil [58] study may have focused on a specific aspect or type of legal question answering, our survey aims to provide a comprehensive overview of the field as a whole, encompassing various approaches and domains. Secondly, our survey employs both quantitative and qualitative analysis techniques to offer a more comprehensive and holistic understanding of the state of the art in legal question answering. Finally, our survey may incorporate more recent literature and developments in the field, as our knowledge cutoff date is more recent compared to Martinez-Gil [58]'s study.

Finally, a broader perspective of AI approaches to the field of legal studies is provided in the recent tutorial presented at ECIR 2023 conference Ganguly et al. [24].¹ Interested readers are encouraged to refer to this resource if they wish to obtain a comprehensive overview of NLP and IR techniques (not necessarily QA) applied on legal documents.

QA methods

We describe now popular methods used for generic and non-domain specific QA systems to provide a necessary background for understanding Legal QA models which will be discussed in Sect. "Legal QA Models".

Question answering (QA) has become an essential tool for extracting information from large amounts of data. Classic machine learning approaches for QA include rule-based systems and information retrieval methods which rely on predefined rules and patterns to match questions with answers. However, these methods lack the ability to understand natural language and adapt to new patterns and changes in the data. On the other hand, modern machine learning approaches such as deep learning and transformer-based models like BERT [18, 69], GPT-2 [70], and GPT-3 [12] leverage advanced algorithms and large amounts of data to train models that can understand natural language and generate accurate responses. These models have been shown to be more effective and robust in handling different language patterns. In this section, we will discuss these approaches in more detail.

Classic machine learning for QA

Rule-based methods: Rule-based methods [31, 77] are a type of classic machine learning approach for QA. They are based on a set of predefined rules and patterns that are used to match questions with answers. These rules are typically created by domain experts or through manual dataset annotation. They are best suited for tasks where the questions and answers can be easily defined using a set of rules, such as in a FAQ [90] system or a medical diagnostic system [75]. However, one of the main limitations of rule-based systems is their lack of ability to understand natural language. They are based on matching keywords or patterns, and they cannot understand the text's meaning. Additionally, these systems can be brittle to changes in the data, as they cannot adapt to new patterns or variations in the language.

Information retrieval (IR) based methods: Information Retrieval (IR) based methods [87] are another classic machine learning approach for QA. These methods rely on pre-processing and indexing the data to make it searchable. They then use algorithms such as cosine similarity [1] or TF-IDF [92] to match the question with the most relevant answer. These methods are best suited for tasks where the questions and answers are already available in a large corpus of text, such as through a search engine [39] or a document retrieval system [16]. However, they are not able to "understand" the meaning of the text and they can provide irrelevant results. These methods are essentially based on matching keywords or patterns, and they are not able to understand the context or

¹ <https://github.com/law-AI/ecir2023tutorial>.

the intent of the question. Additionally, these methods require a large amount of labeled data to work effectively.

Modern machine learning for QA

Deep Learning: Deep learning (DL) is a modern machine learning approach for QA that relies on neural networks to understand natural language. These networks are trained on large amounts of data and are able to understand the meaning of the text. Popular architectures include Recurrent Neural Networks (RNN) Rumelhart et al. [73], Long Short-Term Memory (LSTM) Hochreiter and Schmidhuber [34], and Convolutional Neural Networks (CNN) Krizhevsky et al. [53]. These models can be fine-tuned for specific tasks such as QA.

These models are able to generate accurate responses and adapt to new patterns and changes in the data. They are able to understand the context and the intent of the question, and they can provide relevant and natural-sounding responses. Additionally, these models can be trained on a wide range of tasks such as question answering [3, 76], language translation [29], and text summarization [62, 74, 78].

Transformer-based models: Transformer-based models such as BERT [18] and GPT-2 [70] belong to a type of deep learning approach that has been shown to be very effective in a wide range of natural language processing tasks. These models are based on the transformer architecture, which allows them to learn the context of the text and understand the meaning of the words. The key feature of these models is the use of self-attention mechanisms, which enables them to effectively weigh the importance of different parts of the input when making predictions. This allows for understanding the context of a given question and providing a relevant answer.

BERT is a transformer-based model that was pre-trained on a massive amount of unsupervised data. For the pre-training corpus, BERT used BooksCorpus [98] (800 M words) and English Wikipedia (2,500 M words). The model was trained on a large corpus of unlabelled text data, allowing it to learn the language's general features. BERT is often fine-tuned on a task-specific dataset to perform various natural language understanding tasks such as question answering, sentiment analysis, and named entity recognition. As a pre-trained transformer model, BERT uses a technique called masked language modeling, where certain words in the input are randomly masked, and the model is trained to predict the original word from the context. The second pretraining task is the Next Sentence Prediction which is similar to the Textual Entailment task. BERT is applicable in sentence prediction assignments [81], including text completion and generation [94]. The model has been trained to anticipate a missing word or sequence of words when given context. With its bidirectional design, BERT is able to comprehend contextual information from both the left and right of the target word, rendering it an appropriate choice for sentence prediction tasks where context plays a crucial role in generating accurate results.

GPT-2 is another pre-trained transformer model that is fine-tuned on a task-specific dataset to perform a wide range of natural language understanding and generation tasks, including question answering, text completion, and machine translation. GPT-2 was trained on a massive amount of unsupervised text data, allowing it to generate text similar in style and content to human-written text. Like BERT and other

Table 2 Comparison of classic machine learning and modern transformer-based models

Aspect	Classic machine learning	Modern transformer-based models
Data requirements	Requires large labeled dataset	Can work with a smaller labeled dataset
Feature engineering	Requires manual feature engineering	Automatically learn features from data
Model complexity	Simple models, such as logistic regression or SVM	Complex models, such as BERT, and GPT-2
Training time	Faster	Slower
Accuracy	Lower	Higher
Handling of contextual Information	Limited	Strong ability to handle contextual information
Handling of Unstructured data	Limited	Strong ability to handle unstructured data
Generalization ability	Can generalize well	Can generalize better

transformers, GPT-2 can be fine-tuned on a task-specific dataset to perform a wide range of natural language understanding and generation tasks.

DL models are able to generate more accurate, and natural responses than classic approaches, and they can be used in a wide range of use cases such as question answering [85], language translation [97], and text summarization [56]. They are able to understand the context and the intent of the question, and they can provide relevant and naturally sounding responses.

Table 2 compares classic machine learning and modern transformer-based models in several aspects. In terms of data requirements, classic machine learning models require large labeled datasets for training, whereas modern transformer-based models can work with a smaller amount of labeled data. For feature engineering, classic machine learning models require manual feature engineering, whereas modern transformer-based models can automatically learn features from the data. Classic machine learning models tend to have simple models, such as logistic regression or support vector machines, whereas modern transformer-based models have complex models, such as BERT and GPT-2. Classic machine learning models tend to have faster training time than modern transformer-based models but at the cost of lower accuracy. Modern transformer-based models on the other hand have a strong ability to handle contextual information and unstructured data and better generalization than classic machine learning models.

QA Evaluation Metrics

There are several evaluation metrics commonly used to assess the performance of QA systems. In this section, we discuss some of these metrics and provide the relevant equations.

Accuracy

Accuracy [27] is a simple metric that measures the percentage of correctly answered questions. It is calculated as follows:

$$Accuracy = \frac{\text{number of correctly answered questions}}{\text{total number of questions}} \quad (1)$$

Precision and recall

Precision and recall [27] are two metrics often used in information retrieval tasks and can be applied to QA systems as well. Precision measures the percentage of correct answers among the answers that were provided, while recall measures the percentage of correct answers among all possible correct answers. These metrics can help evaluate how well the system is able to provide accurate answers and identify relevant information. Precision and recall are calculated as follows:

$$Precision = \frac{\text{number of correct answers}}{\text{total number of answers provided}} \quad (2)$$

$$Recall = \frac{\text{number of correct answers}}{\text{total number of possible correct answers}} \quad (3)$$

F1 score

The F1 score [27] is a measure of the system's accuracy that takes both precision and recall into account. It is calculated as the harmonic mean of precision and recall and provides a balanced evaluation of the system's performance. The F1 score is calculated as follows:

$$F1Score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (4)$$

Mean reciprocal rank (MRR)

The MRR [86] is a metric that evaluates the ranking of correct answers. It measures the average of the reciprocal of the rank of the first correct answer, where a higher rank receives a lower score. The MRR is calculated as follows:

$$MRR = \frac{1}{\text{number of questions}} \times \sum_{i=1}^{\text{number of questions}} \frac{1}{\text{rank of first correct answer for question } i} \quad (5)$$

BLEU score

The BLEU score [66] is commonly used in natural language processing tasks, including QA. It measures the similarity between the system's output and the human-generated reference answers based on n-gram matches. It is particularly useful for evaluating the system's ability to generate natural and accurate language. The BLEU score is calculated as follows:

$$BLEU = BP \times \exp \left(\sum_{n=1}^N w_n \log p_n \right) \quad (6)$$

where BP is the brevity penalty, which is used to penalize short system outputs, and p_n is the n-gram precision, which measures the proportion of n-grams in the system output that are also present in the reference answers. The weights w_n are used to give higher importance to higher-order n-grams.

Exact match (EM)

EM measures [96] the percentage of questions that the model answered exactly correctly, without any errors or mistakes. EM is calculated as the ratio of the number of questions for which the model gave an exact match answer to the total number of questions.

Datasets

The following section outlines the top LQA datasets and presents a comprehensive list of open data sources utilized in the reviewed studies, as shown in Table 3. It should be noted that many publicly available data sets are not thoroughly cleaned or preprocessed, making it challenging to assess the effectiveness of various models in future studies. In this section, we highlight 14 LQA datasets, including explaining of their source data, their sizes, and the types of answers provided.

PrivacyQA [71] contains 1,750 questions about the privacy policies of mobile applications. The authors created a corpus of privacy policies collected from 35 mobile applications from the Google Play Store. All the privacy policies are in English and were collected before 1st April 2018. The dataset has been created by crowdsourcing, where users are not given the actual privacy policies but are provided with public information on the Google Play Store. Looking at the information provided, i.e., name, description, and navigable screenshots, crowd workers asked questions about the privacy of user content on that particular application. To answer those questions, seven legal expert annotators were asked to identify the answers from privacy policies. The dataset also categorizes the questions into nine categories such as First party collection/use, Third party sharing/collection, Data Security, Data Retention, User Choice/Control, User Access, Edit and Deletion, Policy Change, International and Specific Audiences, Other.

JEC-QA [95] is a question-answering dataset in the legal domain. The data has been collected from the National Judicial Examination of China and other websites for examinations. JEC-QA contains 26,367 multiple-choice questions along with labels defining the type of questions and the reasoning abilities to answer these questions. The dataset also contains a database of legal knowledge required to answer these exam questions. The database is collected from the National Unified Legal Professional Qualification Examination Counseling book and other Chinese Legal provisions. The book contains 15 topics and 215 chapters with a highly hierarchical form of content. The dataset contains 3,382 different Chinese legal provisions.

The Legal Argument Reasoning Task in Civil Procedure [10] dataset is based on the book the Glannon Guide To Civil Procedure by Glannon, [25] and it is used for a task of legal argument reasoning in civil procedure. It contains multiple choice questions that include a question, answer candidates, a correct answer, a short introduction to the topic of the question and an analysis of why the correct answer is correct. The dataset is split into train, dev and test sets, and the task is defined as identifying whether the answer candidate is correct or incorrect. The authors manually parsed the book and separated the analysis to isolate the relevant aspect for each answer, and this process allowed them to create a binary classification task. The final dataset consists of 918 entries.

French statutory article retrieval dataset (BSARD) [57] is a collection of structured French legal texts, aiming to provide easier access to and analysis of these texts

Table 3 Comparison of legal question answering datasets in terms of languages, source, category, and size

Dataset	Language	Source	Category	Answer type	Size	Year	Question annotators	Availability	Valuation metrics
PrivacyQA ^a	English	Crowd-sourced from Google Play Store	Privacy policies	List of Sentences	1750	2019	Domain experts	Yes	Precision, recall, accuracy
JEC-QA ^b	Chinese	National Judicial Examination of China and other exam websites	Legal knowledge	Single-answer Multiple-answer	26,367	2020	Not mentioned	Yes	Accuracy
BSARD ^c	French	Belgian law articles	Legal texts	Article	1100	2021	Experienced jurists	Yes	Recall, MAP, MRR
PIL	English	Rome I Regulation EC 593/2008, Rome II Regulation EC 864/2007, Brussels I bis Regulation EU 1215/2012	International private law	Articles Recitals Commission Statements	17	2021	Not mentioned	Yes	Recall, precision, F1
COLIEE-2015	Japanese and English	Japanese Legal Bar exams	International private law	Articles Yes/No	412	2015	Not mentioned	Yes	F-measure, Accuracy
VLQA	Vietnamese	Web pages for driver license test	Transportation law domain	Span of words Yes/No	1678	2017	Manually annotated	No	Precision, recall, F1
CUAD ^d	English	Electronic Data Gathering, Analysis, and Retrieval ("EDGAR") system	Legal contract review	Clause	13000	2021	Legal experts	Yes	Precision, recall
Hoppe et al. [35]	German	German Court Rulings, German Legal Documents	Court Rulings	Passage	200	2021	Hand annotated	Yes	Recall, MAP
AILA	Chinese	Chinese Legal Forum	Chinese Law	Span of words	139,468	2020	Professional from law firm	No	MAP, MRR
CJRC	Chinese	Judgments from Supreme Peoples' Court of China	Court Judgments	Span of words yes/no unanswerable	50K	2019	Law experts	No	F1-Score
Collarana et al. [17]	English	MaRisk Regulatory document		Span of words	631	2018	Not mentioned	No	F1-Score, exact match
LCQA	English	Avvo QA forum	User privacy		5,628,689	2022	Not mentioned	Yes	MRR, MAP, precision

Table 3 (continued)

Dataset	Language	Source	Category	Answer type	Size	Year	Question annotators	Availability	Valuation metrics
Kien et al. [42]	Vietnamese	Legal Advice Websites	law and Regulation Documents	Article	5922	2020	Not mentioned	No	Recall@20, NDCG@20
PolicyQA ^e	English	OPP-115 Corpus	Privacy Policies	Span of words	714	2020	Hand annotated	Yes	EM, F1-Score

^a <https://github.com/wasiahmad/PolicyQA>

^b <https://jecqa.thunlp.org/>

^c <https://github.com/maastrichtlawtech/bsard>

^d <https://www.atticusproject.ai.org/cuad>

^e <https://github.com/wasiahmad/PolicyQA>

by researchers, lawyers, and other professionals. To mitigate this problem, experts put together the Belgian Statutory Article Retrieval Dataset (BSARD), which contains over 1,100 legal issues annotated by experienced jurists with pertinent articles from a corpus of over 22,600 Belgian law articles and written entirely in French. Laws, regulations, codes, and other forms of legal writing are included in the dataset. These texts could be useful when one studies criminal law, civil law, or business law in a certain jurisdiction (state or country).

International Private Law (PIL) [80] is a complex legal field with frequently opposing standards between the hierarchy of legal sources, legal jurisdictions, and established procedures. Research on PIL demonstrates the necessity for a link between European and national laws. In this setting, legal professionals must access diverse sources, be able to recall all applicable rules, and synthesize them using case law and interpretation theory concepts. Whenever regulations change frequently or are of sufficient size, this obviously poses a formidable obstacle for people. Automated reasoning over legal texts is not a simple undertaking due to the fact that legal language is highly specialized and in many ways distinct from everyday natural language. This dataset was developed by expanding a previous dataset from Sovrano et al. [79]; it contains questions on the Rome I Regulation EC 593/2008, the Rome II Regulation EC 864/2007, and the Brussels I bis Regulation EU 1215/2012. This legislation has as its sole objective the regulation of matters involving conflicts of law and conflicts of jurisdiction. Therefore, legal specialists were urged to refrain from allowing case law, general principles, or scholarly viewpoints to influence their responses. The objective of this study is to model only the neutral legislative information from the three Regulations, without any other interpretation except the literal one. The incorporation of further information will be left to future studies.

Competition on Legal Information Extraction/Entailment (COLIEE) [45] is a collaborative evaluation task for legal question-answering systems which started in 2014 and which continues to publish a new set of challenges every year since then. Compared to other datasets in legal research, this dataset is the most commonly used by researchers in LQA. The task behind COLIEE aims to establish a standard for evaluating quality assurance (QA) systems in the legal domain and promote research in the field. The challenge is based on a dataset of legal questions and answers provided by the organizers of the COLIEE workshop. The dataset consists of a collection of legal questions, each

with several alternative responses, and the objective of the QA systems is to rank the responses in descending order of importance. COLIEE-2014-COLIEE-2016 is a competition that focuses on legal question-answering tasks containing three subtasks: legal information retrieval, entailment relationship identification, and a combination of both. It is based on a corpus of legal questions culled from Japanese Legal Bar examinations, and the pertinent Japanese Civil Law articles are included. COLIEE (2014–2016) was organized by Juris-informatics (JURISIN) workshop. COLIEE-2017-COLIEE-2022 on the other hand was with the ICAIL conference. In COLIEE-2018, a new corpus based on case Laws of the Federal Court of Canada provided by Compass Law was included.

The purpose of the Vietnamese Legal Question Answering (VLQA) [8] centers on analyzing questions in the legal realm of transportation legislation in Vietnamese. The objective is to extract crucial information such as vehicle type, vehicle action, location, and question type from a legal query expressed in natural language. This information is then utilized to retrieve the answer from a knowledge base. The authors propose a technique that leverages Conditional Random Fields (CRFs) to extract important information from the questions. A corpus of 1,678 questions about Vietnamese transportation law that has been annotated is also presented. The study emphasizes the significance of transportation law in Vietnam due to the prevalence of private automobiles and motorcycles. Simultaneously, a significant number of transportation law infractions have been documented, primarily owing to ignorance or lack of legal knowledge. The authors argue that the existence of a QA system in natural language can be a good option for increasing Vietnamese drivers' awareness and comprehension of transportation law. The paper proposes a method for extracting crucial information from legal questions in the Vietnamese language pertaining to transportation legislation, which can be used to create a question-answering system for this domain.

CUAD [33] is a novel dataset for legal contract review that was compiled with the assistance of dozens of legal professionals from The Atticus Project. It contains over 13,000 expert annotations and more than 500 contracts over 41 label categories. The objective is to highlight the most significant, human-reviewable elements of a contract. CUAD has been annotated by experts for legal contract evaluation and is used to train and assess the performance of NLP models in the legal domain. The dataset includes a collection of legal agreements and annotations submitted by legal professionals. The agreements cover a vast array of legal themes, including contract law, business law, and more. In order to evaluate the performance of models trained on the dataset, the dataset is divided into a training set, a validation set, and a test set. The dataset's annotations include information such as entities, clauses, and clause-to-clause relationships. This information can be used to train models to comprehend the structure and meaning of legal agreements, hence aiding legal contract review activities such as spotting potential flaws and hazards in legal agreements.

Hoppe et al. [35] build an intelligent legal advisor on German Legal documents. They create a data set that consists of 200 hand-annotated question-answer pairs from German Legal documents. As an underlying data source, they use rulings and decisions

of the German court, which are published by openlegalddata.² The documents consist of approximately 200,000 judgments published between 1970 and 2021 from different courts at different levels, such as city, state, and federal. Each document contains meta-data, such as the assigned field of court or law, along with the plain judgment text.

AILA [38] is a question-answering system in the domain of Chinese laws. The LegalQA dataset used in that work, which has 139,468 QA pairs, was developed by gathering QA pairs from a Chinese online legal forum. Real individuals post queries on the forum, and qualified lawyers respond with the appropriate answers. The authors manually construct a KG containing legal concepts and their relations with the assistance of legal professionals for annotation. The legal KG comprises 42,414 legal concepts belonging to 1426 disputes.

CJRC [20] is a Chinese judicial reading comprehension dataset. It comprises 50K question-answer pairs with 10K documents. Underlying documents are judgments published by the Supreme People's Court of China.³ The document collection consists of 5858 criminal documents and 5,737 civil documents. The dataset is created with the assistance of layers by forming four to five question-answer pairs based on the case description. The data set consists of questions with a span of word answers and yes/no, unanswerable questions similar to SQuAD 2.0 and CoQA datasets. The structure is based on a case name containing information like Cause of Action or Criminal Charge, Case Description, and QA Pairs.

Collarana et al. [17] proposed a question-answering dataset based on the MaRisk regulatory documents, which define minimum risk management requirements for banks, insurance, and financial trading companies in Germany. The dataset is created based on the English language MaRisk document. The document is 62 pages long, containing 64 sections and subsections. The dataset comprises of 631 question-answer pairs based on MaRisk document.

LCQA [7] is a large collection of legal questions and answers scraped from the Avvo QA forum. It contains over 5 million questions and has been anonymized to protect user privacy. The questions are organized into categories, with a focus on bankruptcy law in the state of California, and includes 9897 total posts from 3741 lawyers. The dataset also includes relevance labels and query selection to identify experts on a particular category tag, based on engagement filtering and acceptance ratio criteria. This dataset is a valuable resource for researchers and practitioners in the field of legal information retrieval, natural language processing, and legal artificial intelligence.

Kien et al. [42] build a question-answering dataset on Vietnamese legal documents. The authors created two datasets: one is the collection of a legal document corpus containing Vietnamese legal documents. The raw legal documents crawled from official online sites^{4,5}. The raw crawled documents included different versions of each law and regulation. With the lawyers' help, the final corpus contains non-redundant and recent articles, being composed in total of 8586 documents with 117,545 articles. Second is

² <http://openlegalddata.io/>.

³ <https://wenshu.court.gov.cn/>.

⁴ <http://vbpl.vn/tw/pages/home.aspx>.

⁵ <https://thuvienphapluat.vn>.

the question-answer pairs collected from legal advice websites^{6,7,8}. The question dataset contains 5922 queries along with their relevant articles as answers. The question-answer dataset is annotated with lawyers' assistance by mapping current effective articles as answers to questions and removing the old article.

The last Legal QA dataset we discuss, PolicyQA Ahmad et al. [4], is a question-answering dataset based on the privacy policies of 115 websites created using the OPP-115 corpus. The corpus contains 23,000 data practices, 128,000 practice attributes, and 103,000 annotated text spans which the experts manually annotate. The QA dataset contains 714 manually annotated questions based on the privacy policies of 115 websites.

Legal QA models

In recent years, several research studies have been conducted to address the challenge of answering legal questions using natural language processing (NLP) techniques. This section will review some of the most notable studies in this field, highlighting their key contributions and similarities.

Legal information retrieval models

Using a combination of the term frequency-inverse document frequency (tf-idf) model and a support vector machine (SVM) re-ranking model, Kim et al. [49] proposed a system for retrieving legal information and answering yes/no questions. They evaluate their system using a dataset of Japanese civil law articles and questions from legal bar examinations. For information retrieval, the system employs the tf-idf model to retrieve the most pertinent articles for a given query. The SVM re-ranking model is then used to determine the significance of additional features, such as matched phrases between the article and the query. Additionally, lemmatization and dependency parsing are utilized to improve retrieval performance. Experiments revealed that the SVM model containing all three features (lexical words, dependency pairs, and tf-idf scores) performed the best. The system predicts textual entailment for answering yes/no questions by combining word embedding for semantic analysis and paraphrasing for term expansion. They extract features from sentences, such as negation and synonym/antonym relationships, to confirm the correct entailment. In addition to identifying the most pertinent articles and sentences, the system divides them into conditions, conclusions, and exceptions. Experiments revealed that their SVM-based system outperformed previous techniques and achieved an accuracy of 62.14% on the dry run dataset, 55.71% for Phase 2, and 55.79% for Phase 3. When the two phases were combined, the system also performed at its peak.

In Duong and Ho [21], the authors propose vLawyer which is a Vietnamese Question Answering system that uses available Vietnamese resources and tools, such as VietTokenizer, jvnTagger, and Lucene, to build a corpus of legal documents and extract pertinent information to respond to user queries. The system permits users to ask questions in natural language and returns relevant legal articles, clauses, or sentences in response.

⁶ <https://hdpl.moj.gov.vn/Pages/home.aspx>.

⁷ <http://hethongphapluat.com/hoi-dap-phap-luat.html>

⁸ <https://hoidapphapluat.net>.

The system finds answers to user queries using a similarity-based model by extracting candidate passages, constructing a term-document matrix, and calculating similarities using the cosine function and LSI space. The system then combines the results with heuristics to determine which response to return to the user. It selects the first ten results with the highest similarity scores; the more the candidate passage contains the key phrases/keywords of the query, the more likely it is that the passage is the answer. If two candidate passages contain the same number of keyphrases and/or keywords, the system selects the shorter candidate passage. The authors report that their system achieved about 70% precision in legal documents, demonstrating the validity of their approach in the legal document domain.

Kim et al. [44] proposed a QA method for answering yes/no questions on legal bar examinations, which is one of the earliest studies in this field. The authors divided their strategy into two steps: the first relevant legal documents are identified, and in the second step, answers to questions are found through the analysis of relevant documents. The paper focuses on the second task, which is a form of Recognizing Textual Entailment (RTE) in which the input is a question sentence and its corresponding civil law article(s), and the output is a binary answer. A hybrid method is proposed that combines simple rules with an unsupervised learning model employing profound linguistic characteristics. The authors developed a knowledge base for negation and antonym words in the legal domain and employed a two-phase approach to answering yes/no questions. They utilized a dataset of 247 questions paired with civil law articles annotated by legal experts, where 25.63% of questions had a one-to-one correspondence between the question and the article. They evaluated their method using a Korean translation of the original Japanese data, a simple unsupervised learning method, and SVM, a supervised learning model. The outcome demonstrated that the proposed method had an overall accuracy of 61.13%, outperforming both the unsupervised learning technique applied to all questions and the SVM model. In addition, the paper also states that the rule-based approach for simple questions was accurate 68.36% of the time and covered 47.18% of all questions.

A question-answering system for jurisprudential legal questions in the Muslim faith, called KAB, is proposed in Alotaibi et al. [6]. The system utilizes a combination of retrieval-based and generative-based techniques and incorporates prior knowledge sources, such as previous questions, question categories, and Islamic Jurisprudential reference books and sources as a source of context for the produced answer. The architecture of the system includes both generative and retrieval-based methods, where the input question text (Q) is first preprocessed, and then passed to the Knowledge Database, which contains a collection of historical questions, answers, and categories obtained from official online Islamic Jurisprudential legal websites. The KAB system is trained on a dataset of 850,000 entries and evaluated using metrics such as BERTScore and METEOR to measure its performance in terms of precision, recall, and F1 with values of 0.6, 0.4 and 0.48 respectively, and 0.037 for METEOR. The key goal of the proposed system is to provide relevant and high-quality answers to aid in Muslims' daily life decisions, while also reducing the workload on human experts.

Hoppe et al. [36] presented an approach for creating an intelligent legal advisor for German legal documents, using state-of-the-art technologies in the fields of NLP,

semantic search, and knowledge engineering. They have shown that document retrieval and QA are highly relevant problems in the legal field that can be improved by their technology approach, making the work of lawyers more efficient and reducing barriers to society's access to legal information. The authors also described the workflow and underlying technologies in detail and performed experiments on document retrieval in the legal domain. They found that the pre-trained BERT model is not effective out-of-the-box in the legal domain, and performed worse than BM25 in recall and mean average precision (MAP). They also found that dense passage retrieval (DPR) performed better on the GermanQuAD data set, suggesting that fine-tuning the pre-trained model with a transfer learning approach could improve performance. The results show that both BM25 and the pre-trained DPR model are able to retrieve relevant passages for the legal questions. However, BM25 performs better than DPR in terms of recall and MAP on legal documents. When compared to the GermanQuAD data, both BM25 and DPR have a recall score greater than 0.8 and show significantly better scores than the legal data set.

Askari et al. [7] proposed methods for expert finding in the legal community question answering (CQA) domain. The goal is for citizens to find lawyers based on their expertise. The authors define the task of lawyer finding and release a test collection for the task. The authors present two types of baseline models for ranking lawyers: document-level and candidate-level probabilistic language models. These models are based on the set of answers written by a lawyer, which is considered the proof of expertise. They also present a Vanilla BERT model, which is a pre-trained BERT model fine-tuned on their dataset in a pairwise cross-entropy loss setting. The authors also proposed a new method which creates four query-dependent profiles for the lawyers. Each profile consists of text that is sampled to represent different aspects of a lawyer's answers such as comments, sentiment-positive, sentiment-negative, and recency. They combine these query-dependent profiles with existing expert finding methods and show that taking into account different lawyer profile aspects improves the best baseline model.

Hoshino et al. [37] proposed architecture of the question-answering system for legal bar exams consists of two parts: a related article search part and a question-answering part. Both parts use predicate argument structure analysis, which compares a pair of sentences based on the case roles of the arguments. The article search part searches for related articles by searching sentences with the same structure. The question-answering part compares whether each sentence represents the same event. The system utilizes a legal term dictionary and an additional dictionary created from morphological analysis results of past legal bar exam problems. The system also uses a dependency parser to make chunks of morphemes and form clause units, which are used to recognize condition clauses and main clauses. The system outputs results using different modules and selects the final answer using SVM by learning each module's confidence value. The authors suggest that tasks such as conditional sentence extraction, person role extraction, and person relationship extraction are important to improve the overall correct answer rate by 70%.

McElvain et al. [60] introduce a non-factoid question-answering (QA) system for the legal domain, WestSearch Plus. WestSearch Plus aims to provide succinct, one-sentence responses to basic legal questions, regardless of the topic or jurisdiction. Using machine

learning algorithms, gazetteer lookup taggers, statistical taggers, and word embedding models, the system is trained on a large corpus of question-answer pairs to predict parts of speech, NP and VP chunks, syntactic dependency relations, semantic roles, named entities and legal concepts, semantic intent, and alignment between noun phrases, dependency relations, and verb phrases. The initial set of questions is extracted from the query logs of Westlaw⁹ (a legal search engine). The answer corpus comprises approximately 22 million human-written, one-sentence summaries of US court case documents spanning more than a century of case law. The system also employs an ensemble model of weak learners to combine all the features and rank each QA pair independently based on a score representing the candidate's likelihood of being the correct answer. The system also determines whether or not to display a response based on the probability score of that response. According to the authors, the proposed model is evaluated based on the Answered at 3 metric, which measures the proportion of questions with at least one correct answer among the top three responses. The system obtained a 90% Answered at 3 metric for correct responses and a 1.5% Answered at 3 metric for incorrect responses. While determining the thresholds for displaying answers, the company also weighed the system's coverage (the number of user questions for which answers are displayed) against the thresholds.

In Martinez-Gil et al. [59], the authors proposed multiple choice question answering system using reinforced co-occurrence analysis was tested on a dataset of legal questions randomly selected from books from the Oxford University Press. The authors used the accuracy metric to determine the performance of the system. The results show that the system was able to correctly answer 13 out of 20 questions, resulting in an accuracy of 65%. The authors also compare their results to other approaches without machine learning capabilities, noting that their approach is able to offer good results at an affordable cost without the need for training and with a high level of interpretability. However, they acknowledge that the system still faces some obstacles in its development related to the amount of engineering work required to tune the parameters involved in the information retrieval pipeline properly.

Textual entailment models

In Kim et al. [48] describes a legal question-answering system that exploits legal information retrieval and textual entailment using a deep convolutional neural network (CNN). Using training/test data from the Competition on Legal Information Extraction/Entailment (COLIEE), which focuses on answering yes/no questions from Japanese legal bar exams, the system is evaluated. The system is comprised of three phases: ad-hoc legal information retrieval, textual entailment, and a learning model-driven combination of the first two phases. Using a combined TF-IDF and Ranking SVM information retrieval component, in phase 1 the system identifies relevant Japanese civil law articles for a legal bar exam query. In phase 2, the system provides "Yes" or "No" responses to previously unseen queries by comparing the extracted query meanings with relevant articles. The textual entailment component of the system is enhanced by a CNN and

⁹ <https://legal.thomsonreuters.com/en/products/westlaw-edge>

dropout regularization. The results demonstrate that this deep learning-based method outperforms an SVM-based supervised baseline model and K-means clustering. This is the first study to apply deep learning to legal question answering for textual entailment. On the dry run data of COLIEE 2014, the legal question-answering system utilizing a convolutional neural network (CNN) with pre-trained semantic word embeddings and dropout regularization performed the best, according to the study's findings. With an input layer dropout rate of 0.1, a hidden layer dropout rate of 0.6, and 100 hidden layer nodes, the system achieved an accuracy of 63.87%. The results indicate that using dropout regularization improves the system's performance by 1.22% when it was not implemented. The system outperformed an SVM-based model with an accuracy of 60.12%, a model proposed by Kim et al. [43] that utilized linguistic features for SVM learning, and a model that incorporated rule-based method and k-means clustering.

In Kim et al. [50], the authors present a system for answering legal questions that are based on a Siamese convolutional neural network (CNN) for textual entailment. The system is designed to classify legal bar exam questions as "yes" or "no" based on the question's semantic similarity to the corresponding law statutes. The system employs a CNN with convolution, max pooling, and rectified linear unit (ReLU) layers, as well as a fully connected top layer. CNNs are trained with a contrastive loss function that combines the distance between the question and statute vectors and the label (yes or no). The authors preprocess the data by removing stop words and performing stemming, then use a CNN with three layers to extract word features from the question and statute segments. The question and statute vectors are subjected to the convolutional layer, and a max pooling layer is applied on top of the CNN output to extract the highest contributing local features and generate a fixed-length feature vector. To prevent overfitting, the authors employ a technique known as dropout, in which a random number of feature detectors are omitted from each training case. The authors evaluated their system using training data from COLIEE 2014 (dry run) and test data from COLIEE 2015 (formal run) for training and validation, respectively. They examined whether the test data and training data overlapped. The COLIEE 2014 dataset has a balanced distribution of positive and negative samples (55.87% yes, 44.13% no), and the baseline accuracy for the true/false assessment is 55.87% (always returning "yes"). The authors trained on 179 questions from the COLIEE 2014 dry run data and achieved a 64.25% accuracy rate. Notably, the authors have not provided any results comparing their system to other cutting-edge systems for legal information extraction/entailment.

Frame-based models

Describe the legal question-answering system using FrameNet for the COLIEE 2018 shared task in Taniguchi et al. [84]. The task involves determining whether a given text from the Japanese bar examination is true or false. The system employs a FrameNet-based semantic database and a predicate-argument structure analyzer to identify semantic correspondences between problem sentences and knowledge source sentences. The authors apply their frame-based system to the COLIEE 2018 task and compare it to their previous system from COLIEE 2017, discovering that, on average, the frame-based system achieves higher scores. Additionally, they utilize the COLIEE training dataset to evaluate the performance of the system and investigate the effects of frame

information. In the article, FrameNet is a lexical database used to identify semantic correspondences between problem sentences and knowledge source sentences in the legal question-answering system. FrameNet is based on the theory of frame semantics, which postulates that people comprehend the meaning of words based on the images they conjure. The database contains both frames and lexical units (LUs), which are the words that evoke the frames. Frame Elements (FEs) are the semantic roles within the frame and are contained within the frames. The authors use FrameNet to compare pairs of frame candidates and the Dijkstra Algorithm to calculate the confidence between two frames. To determine the similarity between frames, they assign different frame relation types, such as inheritance and using weight values. The value of confidence is computed by multiplying the weights of the frame relations on the path. The authors then compare the clauses of civil law articles and legal bar exams extracted by their rule-based system to answer legal yes/no questions. The results demonstrate that the system is effective, with an average accuracy of approximately 67%, and that frame information is essential for answering legal questions. The authors also experimented with various combinations of modules and threshold values and discovered that the system performed optimally with a threshold of 0.90 and Japanese LUs. They conclude that the system is promising and that there is room for improvement in terms of its precision.

Another legal yes/no question-answering system was developed by Taniguchi and Kano [83] to answer questions regarding the legal domain of a statute. The system utilized case-role analysis to determine the correspondences of roles and relationships between given problem sentences and knowledge source sentences. The system was applied to the JURISIN's COLIEE (Competition on Legal Information Extraction/Entailment) 2016 task and performed better than previous task participants, tying for first place in Phase Two of the current year's task. The experiments focused on Phase Two of the COLIEE 2016 Japanese subtask dataset. The formal run of COLIEE 2016 revealed that the methods tied for first place with iLis7 in Phase Two and placed third in Phase Three. The system's iLis7 method is designed to align structures and words embedded within sentence pairs in order to respond to yes/no questions based on relevant legal articles. The alignment-based method is employed to determine the alignments, which is not simple. Observing the data, the system sorts the yes/no questions into a spectrum from easy to difficult. The system includes two knowledge bases: a negation dictionary and an antonym dictionary. The system employs a rule-based approach to answer simple questions, while machine learning addresses more complex categories by utilizing deeper linguistic data.

Knowledge graph (KG) models

Sovrano et al. [79] presented a solution for extracting and making sense of complex information stored in legal documents written in natural language. The proposed solution comprises four primary steps: KG extraction, Taxonomy construction, Legal Ontology Design Pattern alignment, and KG question answering. KG extraction is accomplished by analyzing the grammatical dependencies of tokens extracted by a dependency parser and identifying noun syntagms (concepts) as potential objects and subjects of the triples to extract. The dependency tree extracts all tokens connecting two distinct target concepts in a sentence, constructing a template from these connecting tokens and

target concepts. Taxonomy Construction is used to properly structure the KG. The KG is organized as a light ontology, with a taxonomy serving as its backbone. This enables efficient abstract querying by identifying a concept's types/classes. The taxonomy construction phase entails constructing one or more taxonomies via Formal Concept Analysis (FCA) by exploiting the hypernyms relationships of the concepts in the Knowledge Base (KG). Legal Ontology Design Pattern Alignment is utilized to enhance the quality of the KG structure by aligning it with recognized legal Ontology Design Patterns (ODPs). The KG extraction is considered a bottom-up approach (from concrete documents to abstract ontologies), whereas the pattern-based design of ontologies is considered a top-down approach (from abstract legal concepts identified by experts to their concretization in the legal documents under examination). The top-down approach is more difficult to implement, whereas the bottom-up approach is prone to errors and duplication, frequently yielding inferior results. To address this issue, the authors propose employing a sort of ontological hinge that connects a bottom-up KG with top-down ODPs to leverage both approaches' advantages. Evaluation is conducted to assess the utility of the resulting Knowledge Graph (KG) in relation to the requirements of the legal user. A team of legal experts selected eight pertinent questions and evaluated the accuracy of the algorithm's responses. The algorithm attained an average top-five recall rate of 34.91%. The results indicate that the QA algorithm is deficient in reasoning, indicating the need for future improvements.

Table 4 presents a summary of various methods, approaches, datasets, key contributions, and accuracy scores of legal question-answering systems. The table includes both factoid and non-factoid QA systems and covers a range of approaches, including hybrid methods, alignment-based approaches, KG extraction, and the use of pre-trained models. From this table, it is evident that there is no single approach that is uniformly successful in the legal domain. The accuracy scores of the different approaches vary widely, with the highest scores being in the 90% range, and the lowest scores being around 60%.

It is also worth noting that some of the approaches in the table use pre-existing datasets, while others create their own datasets. The use of pre-existing datasets, such as COLIEE and PIL, allows for better comparability between different approaches, while the creation of new datasets can be useful in exploring different aspects of legal QA. Regarding key contributions, some of the approaches in the table focus on developing new techniques for analyzing legal texts, such as taxonomic analysis and ontology design pattern alignment. Others focus on leveraging pre-existing resources, such as FrameNet and pre-trained models, to improve accuracy. In general, this table highlights the ongoing challenges in developing accurate legal QA systems. While there have been some notable successes, such as achieving 90% accuracy in some non-factoid QA tasks, there is still a long way to go before fully automated legal question-answering systems become a reality.

Discussion

In this section, we will discuss and summarize the latest trends in legal QA processing and propose some possible extensions while also discussing freely available datasets, evaluation metrics, evaluation tools, and language resources and toolkits. We will begin

Table 4 Comparison of legal question answering methods in terms of approach, dataset, key contributions, and accuracy

Method	Approach	Dataset	Key contributions	Accuracy
Kim et al. [43]	QA for legal bar exams	247 questions	Hybrid method combining simple rules and unsupervised learning using deep linguistic features	61.13%
Taniguchi et al. [83]	Legal yes/no QA system	COLIEE 2016 [45]	Case-role analysis and alignment-based approach for determining alignments	Shared first place in Phase Two, achieved third place in Phase Three
Sovrano et al. [79]	Extracting and making sense of complex information in legal documents	PIL Sovrano et al. [80]	KG extraction, taxonomy construction, legal ontology design pattern alignment, and KG question answering	Top5-recall of 34.91%
McElvain et al. [60]	Non-factoid QA for legal domain	Large corpus of question-answer pairs	Trained on machine learning algorithms, gazetteer lookup taggers, statistical taggers, word embedding models	90% Answered at 3 for correct answers, 1.5% Answered at 3 for incorrect answers
Taniguchi et al. [84]	Legal QA system using FrameNet	COLIEE 2018 [45]	Semantic database based on FrameNet and predicate-argument structure analyzer	67% average accuracy
McElvain et al. [60]	Legal QA system using pre-trained models	22M documents classified to over 120K legal topics	Use of pre-trained models and fine-tuning on legal dataset	90% Answered at 3 for correct answers, 1.5% Answered at 3 for incorrect answers
Kim et al. [48]	Legal QA using CNN	COLIEE [45]	Exploiting legal information retrieval and textual entailment using CNN	63.87%
Kim et al. [48]	Legal information retrieval and QA	Japanese civil law articles and legal bar exams	Combination of tf-idf and SVM re-ranking model, lemmatization and dependency parsing	62.14% on the dry run dataset and 55.71%, 55.79% for Phase 2 and 3 respectively
Duong et al. [21]	Vietnamese QA system	Vietnam's legal documents	Utilization of Vietnamese resources and tools, similarity-based model, and a combination of rule-based and machine learning methods	70% precision
Kim et al. [50]	Textual entailment for legal question answering	COLIEE 2014 training data for training, and the COLIEE 2015 test data for validation	Development of a legal question answering system using Siamese CNNs, preprocessing of data by removing stop words and performing stemming, use of a three-layer CNN to extract word features and a max pooling layer, use of dropout to prevent overfitting	64.25%

Table 4 (continued)

Method	Approach	Dataset	Key contributions	Accuracy
Martinez-Gil et al. [59]	Analyzing co-occurrence patterns in unstructured text corpora	Legal questions randomly selected from books from the Oxford University Press	A new method for the automatic answer of multiple choice questions in the legal domain, ability to reduce workload for professionals in the legal sector, ability to be extrapolated to other specific domains	65%
Hoshino et al. [37]	Predicate argument Structure analysis	COLIEE 2018 [45]	Created legal term dictionary, synonym dictionary for predicates, person estimation feature, four types of question answering modules	70%
Alotaibi et al. [6]	Combination of retrieval-based and generative-based techniques with incorporation of prior knowledge sources such as previous questions, question categories, and Islamic Jurisprudential reference books	Custom dataset	Reduced workload on human experts by providing relevant and high-quality answers to aid in Muslims' daily life decisions.	0.60 precision, 0.40 recall, 0.48 F1 and 0.037 for METEOR
Hoppe et al. [36]	Intelligent legal advisor	German legal documents	Semantic document retrieval and QA using state-of-the-art technologies in NLP, semantic search, and knowledge engineering	0.84 Recall 0.73 MAP

by presenting various legal QA approaches and then delve deeper into the current state of the field.

To gain a better understanding of the current trends in Legal QA methods, we begin by showcasing Fig. 2, which illustrates the number of publication years. The figure reveals a steady rise in the total number of approaches since 2014. Several collaborative methods have been developed to leverage the public's efforts in improving the accuracy of legal QA systems. The latest one was published in 2022. Alotaibi et al. [6] is a Knowledge Augmented BERT2BERT Automated Questions-Answering system for Jurisprudential Legal Opinions. It is a Question-Answering (QA) system based on retrieval augmented generative transformer model for jurisprudential legal questions. The system is designed to solve the problem of jurisprudential legal rules that govern how Muslims react and interact.

The COLIEE competitions held in 2019, 2022, and the upcoming one in 2023 have been instrumental in advancing the field of legal question answering (QA) by providing a standardized platform for evaluating submitted approaches on the same dataset, using the same metrics, and even the same published evaluation tool. The competitions have been running since 2007 and have evolved over time to include a range of subtasks related to legal information extraction and entailment. By participating in these competitions, researchers and practitioners have been able to test and refine their techniques and approaches in a standardized environment, thus paving the way for more effective and accurate legal problem-solving. For better views on the performance of methods on each dataset, we provide a summary Table 4. This table summarizes methods with respect to working or being tested on either public datasets or private ones. This section is important as it provides a clear overview of the different features that have been considered in the model, and helps readers to understand the methodology and approach taken by the authors.

After thoroughly analyzing and studying several research papers in the field of Legal QA, we have identified several common themes and approaches that could be used as guidelines and potential directions for future research. In the following two sub-sections, we recommend some guidelines and potential directions in the two following sub-sections.

Guidelines for legal QA

Based on our analysis of the literature, we recommend that future Legal QA research focus on the following guidelines:

- Use of legal-specific knowledge bases: utilizing legal-specific knowledge can help to improve the accuracy and efficiency of Legal QA systems.
- Incorporation of domain-specific features: incorporating domain-specific features such as legal concepts, entities, and relations can improve the performance of Legal QA systems.
- Development of multi-stage models: developing multi-stage models that incorporate both retrieval and extraction stages can help to improve the accuracy and efficiency of Legal QA systems.

- Data augmentation techniques can be used to artificially expand the size of a given dataset, which can help to improve the performance of machine learning models. In the case of Legal QA, question answer data augmentation involves generating additional training examples by modifying the phrasing or wording of existing questions and answers in the dataset. This approach can help to increase the diversity of the training data and improve the model's ability to handle variations in the wording of questions and answers.

Potential extensions

Along with the guidelines, we suggest some potential directions for developing post-processing approaches.

- While several datasets are commonly utilized to evaluate the performance of various Legal QA approaches, only a limited number of these datasets are freely accessible. These publicly available datasets serve as valuable resources, enabling researchers to compare the effectiveness of their methods and gain a better understanding of their strengths and limitations. However, it should be noted that even when using the same dataset, the manner in which the training, development, and testing data are divided can lead to challenges when attempting to make effective comparisons between different approaches. This highlights the importance of establishing clear and consistent evaluation protocols in Legal QA research, which can help to ensure that results are reproducible and comparable across studies.
- Integration of Explainable AI techniques: One potential extension is to explore the integration of explainable AI technique techniques such as attention visualization and explanation generation. This can help to provide transparency and interpretability to Legal QA systems, enabling users to understand the reasoning behind the system's outputs.
- Development of interactive Legal QA systems: Another potential extension is the development of interactive Legal QA systems that allow users to interact with the system and provide feedback on the accuracy and relevance of the system's outputs. This can help to improve the user experience and enable the system to learn from user feedback.
- Investigation of Legal QA for specific legal domains: While Legal QA has been primarily focused on open-domain question answering, there is a need to investigate Legal QA for specific legal domains such as intellectual property, tax law, and criminal law. This can help to develop domain-specific Legal QA systems that are tailored to the unique requirements and challenges of each domain.
- As the majority of existing approaches in Legal QA are tailored to English language, it is crucial to focus on the development of methods and datasets for Legal QA in other languages.

Conclusion and future work

Legal Question Answering (LQA) is a rapidly growing research field that aims to develop models capable of answering legal questions automatically.

The survey discusses a comprehensive review of recent research on legal question-answering (QA) systems. We highlight the key contributions of these studies, including the development of new taxonomies for legal QA systems, the use of advanced natural language processing (NLP) techniques such as deep learning and semantic analysis, and the incorporation of abundant resources such as legal dictionaries and knowledge bases. The survey also discusses the various challenges that legal QA systems still face and potential directions for future research in this field.

In this survey, several datasets have been detailed, including the Competition on Legal Information Extraction/Entailment (COLIEE), Vietnamese Legal Question Answering (VLQA), Contract Understanding Atticus Dataset (CUAD), Intelligent Legal Advisor on German Legal Documents, AILA, and Chinese Judicial Reading Comprehension (CJRC). Also, in this survey, we discussed different datasets used in legal question-answering studies. The PrivacyQA dataset contains 1,750 questions about the privacy policies of mobile applications collected from 35 mobile applications from the Google Play Store. The JEC-QA dataset contains 26,367 multiple-choice questions in the legal domain, collected from the National Judicial Examination of China and other websites for examinations. The Legal Argument Reasoning Task in Civil Procedure dataset contains 918 multiple-choice questions related to legal argument reasoning in civil procedure. The French statutory article retrieval dataset (BSARD) includes over 1,100 legal issues annotated with pertinent articles from a corpus of over 22,600 Belgian law articles and written entirely in French. Finally, the International Private Law (PIL) dataset contains questions on the Rome I Regulation EC 593/2008, the Rome II Regulation EC 864/2007, and the Brussels I bis Regulation EU 1215/2012, aiming to model only the neutral legislative information from the three regulations, without any other interpretation except the literal one. The COLIEE dataset is a collaborative evaluation task for legal question-answering systems that aims to establish a standard for evaluating quality assurance (QA) systems in the legal domain and promote research in the field.

Finally, the survey further explains that specialized neural network architectures are typically used for QA tasks, such as encoder-decoder models and transformer models that use self-attention mechanisms to capture the meaning and context of the question and generate more accurate answers. The use of these architectures has been the key to the success of deep learning for QA and has enabled the development of highly effective QA systems. The use of NLP techniques in answering legal questions has been the subject of several research studies. Kim et al. [44] proposed a QA method for answering yes/no questions on legal bar examinations. Taniguchi and Kano [83] developed a legal yes/no question-answering system for answering questions regarding the legal domain of a statute. Sovrano et al. [79] presented a solution for extracting and making sense of complex information stored in legal documents written in natural language. McElvain et al. [60] provide one-sentence responses to basic legal questions, regardless of the topic or jurisdiction. Taniguchi et al. [84] described a legal question-answering system using FrameNet for the COLIEE 2018 shared task.

Future work in this field could focus on making less complex models that can answer simple legal questions and ignore different kinds of legal information, like case law and statutes. One area of research that looks good is getting rid of pre-trained language models like BERT and GPT-3 from systems that answer legal questions. Legal question-answering systems may also work less well if they don't use other kinds of knowledge, like legal ontologies and graph-based representations of legal documents. Also, the idea that multimodal data like pictures and videos shouldn't be used in legal question-answering systems might not be a very interesting topic to study. Also, to train and test these systems, we need datasets that are smaller and less varied, as well as evaluation metrics that don't take both precision and recall into account. Lastly, there needs to be less reliable ways to evaluate these systems, like machine evaluation, to make sure they aren't giving wrong or useless answers to legal questions.

Acknowledgements

Not applicable.

Author contributions

Conceptualization, AA, BP and AJ; methodology, AA, BP and AJ; validation AA, BP and AJ; formal analysis, AA, BP and AJ; resources, AA, BP and AJ; data curation, AA, BP and AJ; writing-original draft preparation, AA, BP and AJ; writing-review and editing, AA, BP and AJ; visualization, AA, BP and AJ; supervision, AJ; project administration, AJ. All authors read and approved the final manuscript

Funding

Open access funding provided by University of Innsbruck and Medical University of Innsbruck.

Availability of data and materials

Not applicable.

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Received: 12 April 2023 Accepted: 16 July 2023

Published online: 12 August 2023

References

1. Abbasiantaeb Z, Momtazi S. Text-based question answering from information retrieval and deep neural network perspectives: a survey. *Wiley Interdiscip Rev*. 2021;11(6):e1412.
2. Abdallah A, Hamada M, Nurseitov D. Attention-based fully gated CNN-BGRU for Russian handwritten text. *J Imaging*. 2020;6(12):141.
3. Abdallah A, Kasem M, Hamada MA, Sdeek S. Automated question-answer medical model based on deep learning technology. In: *Proceedings of the 6th International Conference on Engineering & MIS 2020*. 2020. p. 1–8.
4. Ahmad WU, Chi J, Tian Y, Chang K-W. Policyqa: a reading comprehension dataset for privacy policies. *arXiv*. 2020. <https://doi.org/10.48550/arXiv.2010.02557>.
5. Allam AMN, Haggag MH. The question answering systems: a survey. *Int J Res Rev Inform Sci (IJRRIS)*. 2012;2(3).
6. Alotaibi SS, Munshi AA, Farag AT, Rakha OE, Al Sallab AA, Alotaibi MK. Knowledge augmented bert2bert automated questions-answering system for jurisprudential legal opinions. *Int J Comput Sci Netw Secur*. 2022;22(6):346–56. <https://doi.org/10.22937/IJCSNS.2022.22.6.44>.
7. Askari A, Verberne S, Pasi G. Expert finding in legal community question answering. In: *European Conference on Information Retrieval*. Springer. 2022. p. 22–30.
8. Bach NX, Thien THN, Phuong TM, et al. Question analysis for vietnamese legal question answering. In: *2017 9th International Conference on Knowledge and Systems Engineering (KSE)*. IEEE. 2017. p. 154–9.
9. Baral C. *Knowledge representation, reasoning and declarative problem solving*. Cambridge: Cambridge University Press; 2003.
10. Bongard L, Held L, Habernal I. The legal argument reasoning task in civil procedure. *arXiv*. 2022. <https://doi.org/10.48550/arXiv.2211.02950>.

11. Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, et al. Language models are few-shot learners. *Adv Neural Inf Process Syst*. 2020;33:1877–901.
12. Brown TB, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, Agarwal S, Herbert-Voss A, Krueger G, Henighan T, Child R, Ramesh A, Ziegler DM, Wu J, Winter C, Hesse C, Chen M, Sigler E, Litwin M, Gray S, Chess B, Clark J, Berner C, McCandlish S, Radford A, Sutskever I, Amodei D. Language models are few-shot learners. *arXiv*. 2020. <https://doi.org/10.48550/arXiv.2005.14165>.
13. Cadene R, Dancette C, Cord M, Parikh D, et al. Rubi: reducing unimodal biases for visual question answering. *Adv Neural Inf Process Syst*. 2019;32:1564–74.
14. Chen C, Han D, Wang J. Multimodal encoder-decoder attention networks for visual question answering. *IEEE Access*. 2020;8:35662–71.
15. Choi E, He H, Iyyer M, Yatskar M, Yih W-T, Choi Y, Liang P, Zettlemoyer L. Quac: question answering in context. *arXiv*. 2018. <https://doi.org/10.48550/arXiv.1808.07036>.
16. Clarke CL, Terra EL. Passage retrieval vs. document retrieval for factoid question answering. In: *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*. 2003. p. 427–8.
17. Collarana D, Heuss T, Lehmann J, Lytra I, Maheshwari G, Nedelchev R, Schmidt T, Trivedi P. A question answering system on regulatory documents. In: Palmirani M, editor. *Legal knowledge and information systems*. Amsterdam: IOS Press; 2018. p. 41–50.
18. Devlin J, Chang M-W, Lee K, Toutanova K. Bert: pre-training of deep bidirectional transformers for language understanding. *arXiv*. 2018. <https://doi.org/10.48550/arXiv.1810.04805>.
19. Do P-K, Nguyen H-T, Tran C-X, Nguyen M-T, Nguyen M-L. Legal question answering using ranking SVM and deep convolutional neural network. *arXiv*. 2017. <https://doi.org/10.48550/arXiv.1703.05320>.
20. Duan X, Wang B, Wang Z, Ma W, Cui Y, Wu D, Wang S, Liu T, Huo T, Hu Z, et al. Cjrc: a reliable human-annotated benchmark dataset for Chinese judicial reading comprehension. In: Sun M, Huang X, Ji H, Liu Z, Liu Y, editors, et al., *China national conference on Chinese computational linguistics*. Berlin: Springer; 2019. p. 439–51.
21. Duong H-T, Ho B-Q. A Vietnamese question answering system in Vietnam's legal documents. In: *Computer Information Systems and Industrial Management: 13th IFIP TC8 International Conference, CISIM 2014, Ho Chi Minh City, Vietnam, Nov 5–7, 2014*. *Proceedings 14*. Springer. 2014. p. 186–97.
22. Ezzeldin AM, Shaheen M. A survey of arabic question answering: challenges, tasks, approaches, tools, and future trends. In: *Proceedings of The 13th international Arab conference on information technology (ACIT 2012)*. 2012. p. 1–8.
23. Fawei B, Pan JZ, Kollingbaum M, Wyner AZ. A semi-automated ontology construction for legal question answering. *New Gener Comput*. 2019;37(4):453–78.
24. Ganguly D, Conrad JG, Ghosh K, Ghosh S, Goyal P, Bhattacharya P, Nigam SK, Paul S, et al. Legal IR and NLP: the history, challenges, and state-of-the-art. In: Kamps J, et al., editors. *Advances in information retrieval*. Cham: Springer; 2023. p. 331–40.
25. Glannon JW. *Glannon guide to civil procedure: learning civil procedure through multiple-choice questions and analysis*. Boston: Aspen Publishing; 2018.
26. Golub D, He X. Character-level question answering with attention. *arXiv*. 2016. <https://doi.org/10.48550/arXiv.1604.00727>.
27. Goutte C, Gaussier E. A probabilistic interpretation of precision, recall and F-score, with implication for evaluation. In: Losada DE, editor. *Advances in Information Retrieval*. Berlin, Heidelberg: Springer Berlin Heidelberg; 2005. p. 345–59. https://doi.org/10.1007/978-3-540-31865-1_25.
28. Guda V, Sanampudi SK, Manikyamba IL. Approaches for question answering systems. *Int J Eng Sci Technol (IJEST)*. 2011;3(2):990–5.
29. Guo D, Zhou W, Li H, Wang M. Hierarchical LSTM for sign language translation. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 32. 2018.
30. Gupta P, Gupta V. A survey of text question answering techniques. *Int J Comput Appl*. 2012;53(4):1–8.
31. Hayes-Roth F. Rule-based systems. *Commun ACM*. 1985;28(9):921–32.
32. He X, Golub D. Character-level question answering with attention. In: *Proceedings of the 2016 conference on empirical methods in natural language processing*. 2016. p. 1598–607.
33. Hendrycks D, Burns C, Chen A, Ball S. Cuad: an expert-annotated NLP dataset for legal contract review. *arXiv*. 2021. <https://doi.org/10.48550/arXiv.2103.06268>.
34. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput*. 1997;9(8):1735–80.
35. Hoppe C, Pelkmann D, Migenda N, Hötte D, Schenck W. Towards intelligent legal advisors for document retrieval and question-answering in german legal documents. In: 2021 IEEE Fourth International Conference on Artificial Intelligence and Knowledge Engineering (AIKE). IEEE. 2021. p. 29–32.
36. Hoppe C, Pelkmann D, Migenda N, Hötte D, Schenck W. Towards intelligent legal advisors for document retrieval and question-answering in german legal documents. In: 2021 IEEE Fourth International Conference on Artificial Intelligence and Knowledge Engineering (AIKE). IEEE. 2021. p. 29–32. <https://doi.org/10.1109/AIKE52691.2021.00011>.
37. Hoshino R, Taniguchi R, Kiyota N, Kano Y. Question answering system for legal bar examination using predicate argument structure. In: *New Frontiers in Artificial Intelligence: JSAI-isAI 2018 Workshops, JURISIN, AI-Biz, SKL, LENLS, IDAA, Yokohama, Japan, Nov. 12–14, 2018, Revised Selected Papers*. Springer. 2019. p. 207–20.
38. Huang W, Jiang J, Qu Q, Yang M. AILA: a question answering system in the legal domain. In: *IJCAI*. 2020. p. 5258–60.
39. Kadam AD, Joshi SD, Shinde SV, Medhane SP. Notice of removal: question answering search engine short review and road-map to future QA search engine. In: 2015 International Conference on Electrical, Electronics, Signals, Communication and Optimization (EESCO). IEEE. 2015. p. 1–8.
40. Kassner N, Schütze H. BERT-KNN: Adding a kNN search component to pretrained language models for better QA. In: *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics.

2020. p. 3424–30. <https://doi.org/10.18653/v1/2020.findings-emnlp.307>. URL <https://aclanthology.org/2020.findings-emnlp.307>.
41. Khandelwal U, Levy O, Jurafsky D, Zettlemoyer L, Lewis M. Generalization through memorization: nearest neighbor language models. In: International Conference on Learning Representations (ICLR). 2020.
42. Kien PM, Nguyen H-T, Bach NX, Tran V, Le Nguyen M, Phuong TM. Answering legal questions by learning neural attentive text representation. In: Proceedings of the 28th International Conference on Computational Linguistics. 2020. p. 988–98.
43. Kim M, Xu Y, Goebel R. Alberta-kxg: legal question answering using ranking SVM and syntactic/semantic similarity. In: JURISIN Workshop. 2014.
44. Kim M-Y, Xu Y, Goebel R, Satoh K. Answering yes/no questions in legal bar exams. In: JSAI International Symposium on Artificial Intelligence. 2014. p. 199–13.
45. Kim M-Y, Goebel R, Ken S. Coliee-2015: evaluation of legal question answering. In: Ninth International Workshop on Juris-informatics (JURISIN 2015). 2015.
46. Kim M-Y, Xu Y, Goebel R. Applying a convolutional neural network to legal question answering. In: JSAI International Symposium on Artificial Intelligence. Springer. 2015. p. 282–94.
47. Kim M-Y, Xu Y, Goebel R. A convolutional neural network in legal question answering. In: JURISIN Workshop. 2015.
48. Kim M-Y, Xu Y, Goebel R. Applying a convolutional neural network to legal question answering. In: New Frontiers in Artificial Intelligence: JSAI-isAI 2015 Workshops, LENLS, JURISIN, AAA, HAT-MASH, TSDAA, ASD-HR, and SKL, Kanagawa, Japan, Nov 16–18, 2015, Revised Selected Papers. Springer. 2017. p. 282–94.
49. Kim M-Y, Xu Y, Lu Y, Goebel R. Question answering of bar exams by paraphrasing and legal text analysis. In: New Frontiers in Artificial Intelligence: JSAI-isAI 2016 Workshops, LENLS, HAT-MASH, AI-Biz, JURISIN and SKL, Kanagawa, Japan, Nov 14–16, 2016, Revised Selected Papers. Springer. 2017. p. 299–313.
50. Kim M-Y, Lu Y, Goebel R. Textual entailment in legal bar exam question answering using deep siamese networks. In: New Frontiers in Artificial Intelligence: JSAI-isAI Workshops, JURISIN, SKL, AI-Biz, LENLS, AAA, SCIDOCA, kNeXI, Tsukuba, Tokyo, Nov 13–15, 2017, Revised Selected Papers 9. Springer. 2018. p. 35–48.
51. Kolomiyets O, Moens M-F. A survey on question answering technology from an information retrieval perspective. *Inf Sci*. 2011;181(24):5412–34. <https://doi.org/10.1016/j.ins.2011.07.047>.
52. Komeili M, Shuster K, Weston J. Internet-augmented dialogue generation. *arXiv*. 2021. <https://doi.org/10.48550/arXiv.2107.07566>.
53. Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. *Commun ACM*. 2017;60(6):84–90.
54. Lewis P, Denoyer L, Riedel S. Unsupervised question answering by cloze translation. *arXiv*. 2019. <https://doi.org/10.48550/arXiv.1906.04980>.
55. Lin Z, Feng M, Santos CND, Yu M, Xiang B, Zhou B, Bengio Y. A structured self-attentive sentence embedding. *arXiv*. 2017. <https://doi.org/10.48550/arXiv.1703.03130>.
56. Liu Y, Lapata M. Text summarization with pretrained encoders. *arXiv*. 2019. <https://doi.org/10.48550/arXiv.1908.08345>.
57. Louis A, Spanakis G, Van Dijck G. A statutory article retrieval dataset in French. *arXiv*. 2021. <https://doi.org/10.48550/arXiv.2108.11792>.
58. Martinez-Gil J. A survey on legal question answering systems. *arXiv*. 2021. <https://doi.org/10.48550/arXiv.2110.07333>.
59. Martinez-Gil J, Freudenthaler B, Tjoa AM. Multiple choice question answering in the legal domain using reinforced co-occurrence. In: International Conference on Database and Expert Systems Applications. Springer. 2019. p. 138–48.
60. McElvain G, Sanchez G, Matthews S, Teo D, Pompili F, Custis T. Westsearch plus: A non-factoid question-answering system for the legal domain. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval SIGIR19. New York: Association for Computing Machinery. p. 1361–4. 2019. <https://doi.org/10.1145/3331184.3331397>
61. Morimoto A, Kubo D, Sato M, Shindo H, Matsumoto Y. Legal question answering system using neural attention. *COLIEE@ICAIL*. 2017;47:79–89.
62. Mukherjee S, Jangra A, Saha S, Jatowt A. Topic-aware multimodal summarization. In: Findings of the Association for Computational Linguistics: AACL-IJCNLP. 2022. p. 387–98.
63. Nguyen BD, Do T-T, Nguyen BX, Do T, Tjiputra E, Tran QD. Overcoming data limitation in medical visual question answering. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer. 2019. p. 522–30.
64. Nie Y-P, Han Y, Huang J-M, Jiao B, Li A-P. Attention-based encoder-decoder model for answer selection in question answering. *Front Inf Technol Electron Eng*. 2017;18(4):535–44.
65. Pal A, Harper FM, Konstan JA. Exploring question selection bias to identify experts and potential experts in community question answering. *ACM Trans Inf Syst (TOIS)*. 2012;30(2):1–28.
66. Papineni K, Roukos S, Ward T, Zhu W-J. Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics. Philadelphia: Association for Computational Linguistics. 2002. p. 311–8. <https://doi.org/10.3115/1073083.1073135>.
67. Perez E, Lewis P, Yih W-T, Cho K, Kiela D. Unsupervised question decomposition for question answering. *arXiv*. 2020. <https://doi.org/10.48550/arXiv.2002.09758>.
68. Pouyanfar S, Sadiq S, Yan Y, Tian H, Tao Y, Reyes MP, Shyu M-L, Chen S-C, Iyengar SS. A survey on deep learning: Algorithms, techniques, and applications. *ACM Comput Surv*. 2018. <https://doi.org/10.1145/3234150>.
69. Qu C, Yang L, Qiu M, Croft WB, Zhang Y, Iyyer M. Bert with history answer embedding for conversational question answering. In: Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval. 2019. p. 1133–6.
70. Radford A, Wu J, Child R, Luan D, Amodei D, Sutskever I. Language models are unsupervised multitask learners. 2019.

71. Ravichander A, Black AW, Wilson S, Norton T, Sadeh N. Question answering for privacy policies: combining computational and legal perspectives. 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: Association for Computational Linguistics. 2019. p. 4947–58. <https://doi.org/10.18653/v1/D19-1500>.
72. Reddy S, Chen D, Manning CD. CoQA: a conversational question answering challenge. *Trans Assoc Comput Linguist*. 2019;7:249–66.
73. Rumelhart DE, Hinton GE, Williams RJ. Learning internal representations by error propagation. In: Technical report, California Univ San Diego La Jolla Inst for Cognitive Science; 1985.
74. Salchner MF, Jatowt A. A survey of automatic text summarization using graph neural networks. In: Proceedings of the 29th International Conference on Computational Linguistics. Gyeongju: International Committee on Computational Linguistics; 2022. p. 6139–50. URL <https://aclanthology.org/2022.coling-1.536>.
75. Sarrouiti M, Lachkar A, Ouatik SEA. Biomedical question types classification using syntactic and rule based approach. In: 2015 7th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K). vol 1. IEEE. 2015. p. 265–72.
76. Saxe A, Nelli S, Summerfield C. If deep learning is the answer, what is the question? *Nat Rev Neurosci*. 2021;22(1):55–67.
77. Smith S, Kandel A. Verification and validation of rule-based expert systems. Boca Raton: CRC Press; 2018.
78. Song S, Huang H, Ruan T. Abstractive text summarization using LSTM-CNN based deep learning. *Multimed Tools Appl*. 2019;78(1):857–75.
79. Sovrano F, Palmirani M, Vitali F. Legal knowledge extraction for knowledge graph based question-answering. In: Francesconi E, Borges G, Sorge C, editors. Legal knowledge and information systems. Amsterdam: IOS Press; 2020. p. 143–53.
80. Sovrano F, Palmirani M, Distefano B, Sapienza S, Vitali F. A dataset for evaluating legal question answering on private international law. In: Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law. 2021. p. 230–34.
81. Sun Y, Zheng Y, Hao C, Qiu H. NSP-BERT: a prompt-based zero-shot learner through an original pre-training task-next sentence prediction. *arXiv*. 2021. <https://doi.org/10.48550/arXiv.2109.03564>.
82. Talmor A, Herzig J, Lourie N, Berant J. Commonsenseqa: a question answering challenge targeting commonsense knowledge. *arXiv*. 2018. <https://doi.org/10.48550/arXiv.1811.00937>.
83. Taniguchi R, Kano Y. Legal yes/no question answering system using case-role analysis. In: JSAI International Symposium on Artificial Intelligence. Springer. 2016. p. 284–98.
84. Taniguchi R, Hoshino R, Kano Y. Legal question answering system using framenet. In: New Frontiers in Artificial Intelligence: JSAI-isAI 2018 Workshops, JURISIN, AI-Biz, SKL, LENLS, IDAA, Yokohama, Japan, Nov 12–14, 2018, Revised Selected Papers. Springer. 2019. p. 193–206.
85. Van Aken B, Winter B, Löser A, Gers FA. How does bert answer questions? a layer-wise analysis of transformer representations. In: Proceedings of the 28th ACM International Conference on Information and Knowledge Management. 2019. p. 1823–32.
86. Voorhees EM, Tice DM. Building a question answering test collection. In: Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '00. New York: Association for Computing Machinery. 2000. p. 200–7. <https://doi.org/10.1145/345508.345577>.
87. Voorhees EM, Harman DK, et al. TREC: experiment and evaluation in information retrieval, vol. 63. Princeton: Citeseer; 2005.
88. Wang Z, Ng P, Ma X, Nallapati R, Xiang B. Multi-passage BERT: a globally normalized BERT model for open-domain question answering. *arXiv*. 2019. <https://doi.org/10.48550/arXiv.1908.08167>.
89. Weston J, Bordes A, Chopra S, Rush AM, Van Merriënboer B, Joulin A, Mikolov T. Towards AI-complete question answering: a set of prerequisite toy tasks. *arXiv*. 2015. <https://doi.org/10.48550/arXiv.1502.05698>.
90. Xie R, Lu Y, Lin F, Lin L. FAQ-based question answering via knowledge anchors. In: CCF International Conference on Natural Language Processing and Chinese Computing. Springer, 2020. p. 3–15.
91. Yang L, Hu J, Qiu M, Qu C, Gao J, Croft WB, Liu X, Shen Y, Liu J. A hybrid retrieval-generation neural conversation model. In: Proceedings of the 28th ACM international conference on information and knowledge management. 2019. p. 1341–50.
92. Yin Y, Zhang Y, Liu X, Zhang Y, Xing C, Chen H. Healthqa: A Chinese QA summary system for smart health. In: International Conference on Smart Health. Springer. 2014. p. 51–62.
93. Zamani H, Diaz F, Dehghani M, Metzler D, Bendersky M. Retrieval-enhanced machine learning. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '22. New York: Association for Computing Machinery. 2022. p. 2875–86. <https://doi.org/10.1145/3477495.3531722>.
94. Zhang T, Kishore V, Wu F, Weinberger KQ, Artzi Y. BERTscore: evaluating text generation with BERT. *arXiv*. 2019. <https://doi.org/10.48550/arXiv.1904.09675>.
95. Zhong H, Xiao C, Tu C, Zhang T, Liu Z, Sun M. JEC-QA: A legal-domain question answering dataset. In: Proceedings of AAAI. 2020.
96. Zhu F, Lei W, Wang C, Zheng J, Poria S, Chua T-S. Retrieving and reading: a comprehensive survey on open-domain question answering. *arXiv*. 2021. <https://doi.org/10.48550/arXiv.2101.00774>.
97. Zhu J, Xia Y, Wu L, He D, Qin T, Zhou W, Li H, Liu T-Y. Incorporating BERT into neural machine translation. *arXiv*. 2020. <https://doi.org/10.48550/arXiv.2002.06823>.
98. Zhu Y, Kiros R, Zemel R, Salakhutdinov R, Urtasun R, Torralba A, Fidler S. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In: Proceedings of the IEEE international conference on computer vision. 2015. p. 19–27.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.