SYSTEM AND METHOD FOR ANALYZING AND MODELING CONTENT

CROSS-REFERENCE TO RELATED APPLICATIONS

[001] This application claims the benefit of U.S. Application No. 15/995,984, filed June 1, 2018, which is hereby incorporated by reference in its entirety.

TECHNICAL FIELD

[002] The present disclosure relates to parsing and analyzing document text. More specifically, the present disclosure relates to extracting content from aggregated text and analyzing the text to generate logical models.

BACKGROUND

[003] The regulatory landscape is complex and rapidly evolving. Successful trading models require a complete understanding of the global rules of exchange, and compliant trading requires systematic, consistent, and auditable answers. The industry, however, faces numerous challenges. Due to the nature and timing of regulatory rule making, financial institutions have historically built tactical internal solutions, often segregated by asset class and/or by regime. Typically, such solutions operate post-execution, are expensive to maintain, do not allow for full traceability, and cannot offer a single view of the full global cross-regulatory implications of a trade. Without a single global eligibility and obligations decision layer that operates across asset class, it is exceedingly difficult, if not impossible, to manage a sustainable, consistent, auditable, and up-to-date approach.

SUMMARY

[003a] In a first aspect, the present disclosure provides a method for resolving legal or regulatory inquires and generating audit records, the method comprising: aggregating, by a processing device, documents from a plurality of data sources and storing the aggregated documents in a document database, wherein each document comprises legal and regulatory content for a legal and regulatory domain; selecting, by the processing device, a first document from the document database, the first document being expressed in a first format; converting, by the processing device, the first document into a structured data object, wherein the structured data object is expressed in a structured data format independent of the first format of the first document; generating, by the processing device, a plurality of annotations associated with the structured data object, wherein each annotation is associated with a data structure representing logic within the legal and regulatory content, and the plurality of annotations form a comprehensive data structure representing the entire logic within the legal and regulatory

content; storing, by the processing device, the generated annotations in an annotation database comprising annotations associated with the documents in the document database; transforming, by the processing device, the structured data object into a domain-specific data structure, wherein the domain-specific data structure comprises a compiled executable logical model for decision making to resolve transactional inquiries in the legal and regulatory domain, wherein the domain-specific data structure is derived at least partially from one or more logical relationships between the generated annotations and one or more annotations of the annotation database; receiving, by the processing device, a request to resolve an inquiry in the legal and regulatory domain pertaining to a proposed transaction, wherein the request comprises a plurality of transaction-specific parameters; executing, by the processing device, the compiled executable logical model, using the plurality of transaction-specific parameters as inputs, to resolve the inquiry; and generating, by the processing device, an audit record comprising a visualization of the executable logical model, how each decision in the executable logical model was made, and the resolution of the inquiry.

[003b] In a second aspect, the present disclosure provides a system for resolving legal or regulator inquires and generating audit records, the system comprising: a data store for maintaining a document database and an annotation database; and one or more processing devices communicatively coupled to the data store, wherein the one or more processing devices are configured to: aggregate documents from a plurality of data sources and store the aggregated documents in the document database, wherein each document comprises legal and regulatory content for a legal and regulatory domain; select a first document from the document database, the first document being expressed in a first format; convert the first document into a structured data object, wherein the structured data object is expressed in a structured data format independent of the first format of the first document; generate a plurality of annotations associated with the structured data object, wherein each annotation is associated with a data structure representing logic within the legal and regulatory content, and the plurality of annotations form a comprehensive data structure representing the entire logic within the legal and regulatory content; store the generated annotation in the annotation database comprising annotations associated with the documents in the document database; and transform the structured data object into a domain-specific data structure, wherein the domain-specific data structure comprises a compiled executable logical model for decision making to resolve transactional inquiries in the legal and regulatory domain, wherein the domain-specific data structure is derived at least partially from one or more logical relationships between the generated annotations and one or more annotations of the annotation database; receive a request

to resolve an inquiry in the legal and regulatory domain pertaining to a proposed transaction, wherein the request comprises a plurality of transaction-specific parameters; execute the compiled executable logical model, using the plurality of transaction-specific parameters as inputs, to resolve the inquiry; and generate an audit record comprising a visualization of the executable logical model, how each decision in the executable logical model was made, and the resolution of the inquiry.

[003c] In a third aspect, the present disclolsure provides a non-transitory computer-readable storage medium having instructions encoded thereon that, when executed by a processing device, cause the processing device to resolve legal or regulatory inquires and generate audit records by performing at least: aggregate documents from a plurality of data sources and store the aggregated documents in a document database, wherein each document comprises legal and regulatory content for a legal and regulatory domain; select a first document from the document database, the first document being expressed in a first format; convert the first document into a structured data object,wherein the structured data object is expressed in a structured data format independent of the first format of the first document; generate a plurality of annotations associated with the structured data object, wherein each annotation is associated with a data structure representing logic within the legal and regulatory content, and the plurality of annotations form a comprehensive data structure representing the entire logic within the legal and regulatory content; store the generated annotation in an annotation database comprising annotations associated with the documents in the document database; and transform the structured data object into a domain-specific data structure, wherein the domain-specific data structure comprises a compiled executable logical model for decision making to resolve transactional inquiries in the legal and regulatory domain, wherein the domain-specific data structure is derived at least partially from one or more logical relationships between the generated annotation and one or more annotations of the annotation database; receive a request to resolve an inquiry in the legal and regulatory domain pertaining to a proposed transaction, wherein the request comprises a plurality of transaction-specific parameters; execute the compiled executable logical model, using the plurality of transaction-specific parameters as inputs, to resolve the inquiry; and generate an audit record comprising a visualization of the executable logical model, how each decision in the executable logical model was made, and the resolution of the inquiry.

BRIEF DESCRIPTION OF THE DRAWINGS

[004] In order to facilitate a fuller understanding of the present disclosure, reference is now

made to the accompanying drawings, in which like elements are referenced with like numerals. These drawings should not be construed as limiting the present disclosure, but are intended to be exemplary only.

[005] FIG. 1 illustrates an exemplary system architecture in accordance with an embodiment of the present disclosure.

[006] FIG. 2 illustrates an exemplary data flow model in accordance with an embodiment of the present disclosure.

[007] FIG. 3A is a flow diagram illustrating a method for aggregating, parsing, and annotating regulatory content in accordance with an embodiment of the present disclosure.

CLAIMS

WHAT IS CLAIMED IS:

1. A method for resolving legal or regulatory inquires and generating audit records, the method comprising:

aggregating, by a processing device, documents from a plurality of data sources and storing the aggregated documents in a document database, wherein each document comprises legal and regulatory content for a legal and regulatory domain;

selecting, by the processing device, a first document from the document database, the first document being expressed in a first format;

converting, by the processing device, the first document into a structured data object, wherein the structured data object is expressed in a structured data format independent of the first format of the first document;

generating, by the processing device, a plurality of annotations associated with the structured data object, wherein each annotation is associated with a data structure representing logic within the legal and regulatory content, and the plurality of annotations form a comprehensive data structure representing the entire logic within the legal and regulatory content;

storing, by the processing device, the generated annotations in an annotation database comprising annotations associated with the documents in the document database;

transforming, by the processing device, the structured data object into a domain-specific data structure, wherein the domain-specific data structure comprises a compiled executable logical model for decision making to resolve transactional inquiries in the legal and regulatory domain, wherein the domain-specific data structure is derived at least partially from one or more logical relationships between the generated annotations and one or more annotations of the annotation database;

receiving, by the processing device, a request to resolve an inquiry in the legal and regulatory domain pertaining to a proposed transaction, wherein the request comprises a plurality of transaction-specific parameters;

executing, by the processing device, the compiled executable logical model, using the plurality of transaction-specific parameters as inputs, to resolve the inquiry; and

generating, by the processing device, an audit record comprising a visualization of the executable logical model, how each decision in the executable logical model was made, and the resolution of the inquiry.

2. The method of claim 1, wherein one or more of the annotations comprises a taxonomy of a rule, a scope to which the rule applies, and a classification of the annotation.

3. The method of claim 2, wherein the one or more of the annotations further comprises one or more of a definition of a legislative term, a definition of a data element, a cross-reference to another annotation, an expression of a relationship between entities, or an expression of a relationship between data elements.

4. The method of claim 1, wherein generating the plurality of annotations comprises extracting a rule-specific indicator from the legal and regulatory content.

5. The method of claim 4, wherein the processing device utilizes a natural language processing algorithm to identify the rule-specific indicator within the legal and regulatory content.

6. The method of claim 1, wherein the request to resolve an inquiry in the legal and regulatory domain pertaining to a proposed transaction comprises determining obligations of one or more parties associated with the transaction; and wherein the method further comprises generating, by the processing device, an indication of the obligations of the one or more parties based on the domain-specific data structure using the plurality of transaction-specific parameters as inputs.

7. The method of claim 1, further comprising determining, by the processing device, that the first document corresponds to an updated version of a previously-stored document in the document database, wherein the updated version corresponds to legal and regulatory content not present in the previously stored document.

8. The method of claim 1, wherein aggregating the documents from the plurality of data sources comprises: identifying, by the processing device, one or more documents available from the plurality of data sources that have associated hash values that differ from hash values associated with any of the documents previously stored in the document database; and retrieving, by the processing device, the one or more documents and storing the one or more documents in the document database.

9. The method of claim 1, further comprising selecting, by the processing device, the plurality of data sources from which to aggregate the documents based on a target regulation category.

10. The method of claim 1, further comprising periodically reaggregating, by the processing device, the documents from the plurality of data sources and calculating a hash value for each document to detect changes to the documents.

11. The method of claim 1, wherein the legal and regulatory content comprises one or more of: a law, a regulation, a rule, a legal text, a policy text, and a legal guidance.

12. The method of claim 1, wherein the visualization comprises a logic flow chart.

13. A system for resolving legal or regulator inquires and generating audit records, the system comprising:

a data store for maintaining a document database and an annotation database; and

one or more processing devices communicatively coupled to the data store, wherein the one or more processing devices are configured to:

aggregate documents from a plurality of data sources and store the aggregated documents in the document database, wherein each document comprises legal and regulatory content for a legal and regulatory domain;

select a first document from the document database, the first document being expressed in a first format;

convert the first document into a structured data object, wherein

the structured data

object is expressed in a structured data format independent of the first format of the first document;

generate a plurality of annotations associated with the structured data object, wherein each annotation is associated with a data structure representing logic within the legal and regulatory content, and the plurality of annotations form a comprehensive data structure representing the entire logic within the legal and regulatory content;

store the generated annotation in the annotation database comprising annotations associated with the documents in the document database; and

transform the structured data object into a domain-specific data structure, wherein the domain-specific data structure comprises a compiled executable logical model for decision making to resolve transactional inquiries in the legal and regulatory domain, wherein the domain-specific data structure is derived at least partially from one or more logical relationships between the generated annotations and one or more annotations of the annotation database;

receive a request to resolve an inquiry in the legal and regulatory domain pertaining to a proposed transaction, wherein the request comprises a plurality of transaction-specific parameters;

execute the compiled executable logical model, using the plurality of transaction-specific parameters as inputs, to resolve the inquiry; and

generate an audit record comprising a visualization of the executable logical model, how each decision in the executable logical model was made, and the resolution of the inquiry.

14. The system of claim 13, wherein one or more of the annotations comprises a taxonomy of a rule, a scope to which the rule applies, and a classification of the annotation, and wherein the one or more of the annotations further comprises one or more of a definition of a legislative term, a definition of a data element, a cross-reference to another annotation, an expression of a relationship between entities, or an expression of a relationship between data elements.

15. The system of claim 13, wherein to generate the plurality of annotations, the processing device is further configured to extract a rule-specific indicator from the legal and regulatory content, and wherein the processing device is further configured to utilize a natural language processing algorithm to identify the rule-specific indicator within the legal and regulatory content.

16. The system of claim 13, wherein the request to resolve an inquiry in the legal and regulatory domain pertaining to a proposed transaction comprises determining obligations of one or more parties associated with the transaction, and wherein the one or more processing devices are further configured to generate an indication of the obligations of the one or more parties based on the domain-specific data structure using the plurality of transaction-specific parameters as inputs.

17. The system of claim 13, wherein the one or more processing devices are further configured to periodically reaggregate the documents from the plurality of data sources and calculate a hash value for each document to detect changes to the documents.

18. The system of claim 13, wherein the legal and regulatory content comprises one or more of: a law, a regulation, a rule, a legal text, a policy text, and a legal guidance.

19. The system of claim 13, wherein the visualization comprises a logic flow chart.

20. A non-transitory computer-readable storage medium having instructions encoded thereon that, when executed by a processing device, cause the processing device to resolve legal or regulatory inquires and generate audit records by performing at least:

aggregate documents from a plurality of data sources and store the aggregated documents in a document database, wherein each document comprises legal and regulatory content for a legal and regulatory domain;

select a first document from the document database, the first document being expressed in a first format;

convert the first document into a structured data object,wherein the structured data object is expressed in a structured data format independent of the first format of the first document;

generate a plurality of annotations associated with the structured data object, wherein each annotation is associated with a data structure representing logic within the legal and regulatory content, and the plurality of annotations form a comprehensive data structure representing the entire logic within the legal and regulatory content;

store the generated annotation in an annotation database comprising annotations associated with the documents in the document database; and

transform the structured data object into a domain-specific data structure, wherein the domain-specific data structure comprises a compiled executable logical model for decision making to resolve transactional inquiries in the legal and regulatory domain, wherein the domain-specific data structure is derived at least partially from one or more logical relationships between the generated annotation and one or more annotations of the annotation database;

receive a request to resolve an inquiry in the legal and regulatory domain

pertaining to a proposed transaction, wherein the request comprises a plurality of transaction-specific parameters;

execute the compiled executable logical model, using the plurality of transaction-specific parameters as inputs, to resolve the inquiry; and

generate an audit record comprising a visualization of the executable logical model, how each decision in the executable logical model was made, and the resolution of the inquiry.

21. The non-transitory computer-readable storage medium of claim 20, wherein one or more of the annotations comprises a taxonomy of a rule, a scope to which the rule applies, and a classification of the annotation, and wherein the one or more of the annotations further comprises one or more of a definition of a legislative term, a definition of a data element, a cross-reference to another annotation, an expression of a relationship between entities, or an expression of a relationship between data elements.

22. The non-transitory computer-readable storage medium of claim 20, wherein to generate the plurality of annotations, the instructions, when executed by the processing device, further cause the processing device to extract a rule-

specific indicator from the legal and regulatory content, and wherein the instructions, when executed by the processing device, further cause the processing device is to utilize a natural language processing algorithm to identify the rule-specific indicator within the legal and regulatory content.

23. The non-transitory computer-readable storage medium of claim 20, wherein the request to resolve an inquiry in the legal and regulatory domain pertaining to a proposed transaction comprises determining obligations of one or more parties associated with the transaction; and wherein the instructions, when executed by the processing device, further cause the processing device to generate an indication of the obligations of the one or more parties based on the domain-specific data structure using the plurality of transaction-specific parameters as inputs.

24. The non-transitory computer-readable storage medium of claim 20, wherein the instructions, when executed by the processing device, further cause the processing device to periodically reaggregate the documents from the plurality of data sources and calculate a hash value for each document to detect changes to the documents.

25. The non-transitory computer-readable storage medium of claim 20, wherein the legal and regulatory content comprises one or more of: a law, a regulation, a rule, a legal text, a policy text, and a legal guidance.

26. The non-transitory computer-readable storage medium of claim 20, wherein the visualization comprises a logic flow chart.