JANUARY / FEBRUARY 2020

Strategies

THE JOURNAL OF LEGAL MARKETING

2020 VISION

The Technology and Trends Driving Us Forward

+

Using Knowledge Graphs

ABA Model Rule Updates

Merger Mania

LEGAL MARKETING ASSOCIATION
LMA®
THE AUTHORITY FOR LEGAL MARKETING

USING KNOWLEDGE GRA

TO DELIVER BETTER AND FASTER INS

BY DAVID KAMIEN

APHS
IGHTS

Which lawyers speak Spanish, are veterans and have worked on pharma deals over $1 billion in the last three years? This is the type of question law firm leaders may expect their business development (BD) or marketing teams to quickly answer. Unfortunately, most firms must use manual methods to gather this information to inform their BD efforts.

While technology is often used to make legal work more efficient, firms can also utilize advanced analytics software to similarly improve their BD efforts.

The next generation of law firm marketing and BD applications will be powered by knowledge graphs (KGs), which enable law firms to more efficiently integrate data and quickly deliver important insights to the right people. Currently, standard online searches return documents and web pages as results, which users must sift through to find answers. That takes time. It would be more beneficial if queries were answered by returning specific information, which KGs can do.

Interestingly enough, we use KGs every day. For example, you use the Google KG when, in response to typing, "Who directed 'A Star Is Born'?" you see an answer in a box (or knowledge panel) with information about Bradley Cooper above a page of results, streamlining your search.

Systems powered by knowledge graphs will help firms spot opportunities sooner, author higher-value client alerts, foster cross-selling and respond to RFPs. New systems will proactively analyze the relationships between markets, clients, matters, lawyers and practices to spot growth opportunities.

This article explains KGs, how they will power next-generation marketing and BD software systems, and provides use cases that leverage KGs to help law firms compete successfully and win more work.

Technology Background: Ontologies

To understand how a KG works, one must understand the concept of an ontology. In computer and information sciences, an ontology classifies information on a subject into a set of concepts and categories to show their properties and the relations between them. For example, all of human anatomy is organized in The Foundational Model of Anatomy, which is an ontology itself.

An ontology provides an abstract data model (or schema) that states both the definitions of terms and the relationships among those terms in a standardized way that is machine-readable. The concepts, or entities, in the ontology can represent objects, people, places and organizations, among other things, in a specific real-world domain. The ontology defines possible types and attributes or the interrelationships between entities. With an ontology that provides machine-understandable definitions of what the concept "is a kind of" or "located in," software can apply machine reasoning and use inference to deduce new conclusions from relations.

Technology Background: Knowledge Graphs

KGs are repositories that store data and represent knowledge about entities, relations and their abstractions in a machine understandable way.

An ontology isn't a KG until it is connected to the things it describes, such as entity instances (e.g., specific movies and directors), database records or document metadata. In other words, a KG combines both:

(1) a formally-structured, machine-intelligible, knowledge organization system (such as an ontology, a taxonomy or a thesaurus), that defines and represents entities and the relationships between them; and (2) metadata about content, both stored in a graph database capable of supporting various types of sophisticated queries. The creator of the KG declares, in code, that the data points in the graph have the meanings, types and properties defined in the applicable ontologies. A KG will generally have one or more ontologies covering different domains and subdomains, alongside the data that follows those ontologies.

KGs integrate data from internal and external sources with different low-level data schemas. A KG can contain data from text documents, databases or other knowledge bases, and it can be obtained from scanning a database and analyzing files, directories or activity logs. KGs integrate the data into graph-based structures with nodes and directed edges, or links.

Imagine a series of bubbles (called nodes) representing people or companies, connected by lines (called edges) that link the nodes to form a subject-predicate-object group (referred to as a "triple") between the nodes. With nodes and links, you can graphically express things like "Person, is a kind of, Mammal" or "Wall St., is located in, Manhattan." (See image below.)

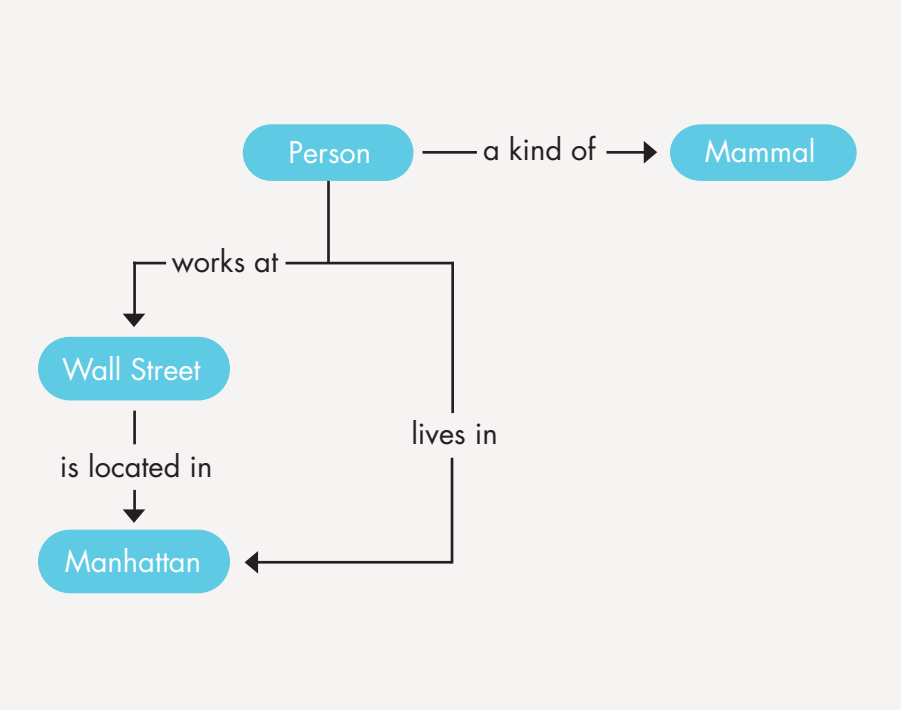

While ontologies are generally hand-curated and contain smaller collections of assertions, KGs are created with software and can therefore include billions of assertions that serve as repositories of highly findable data that is reusable for various use cases.

While an ontology is schema-oriented and a knowledge base is fact-oriented, in the KG, you can have both a schema and a set of facts ("A relation B," or, "Paris isA City," "Paris hasInhabitants 2M," etc.). That is what enables a KG to easily answer to questions like "Number of inhabitants in Paris?"

Apart from ontology-driven inferencing, the KG enables a firm to represent and query, in a natural way, complex relationship patterns, transitive relations and other property paths, deeply nested and irregular tree-like structures, and many other naturally occurring data patterns that do not easily fit in a tabular format. So querying for complex patterns allows us to ask the simple question, "How are A and B related?" With a KG (or a pure graph database), you'd find the path in the network of relationships between A and B, so you could discover that A is on the board with C, at company X, whose CEO is B.

KGs are typically built on top of specialized graph databases, which are well-suited for storing information about relationships among entities, for accessing diverse types of information because they store semantics in a structured format, making it easy for other systems to use and integrate the data.

KGs built on top of Resource Description Framework (RDF) databases support effortless merging of datasets, external cross-linking, and knowledge reuse (see w3.org/DesignIssues/LinkedData.html and https://5stardata.info/en/). RDF is a World Wide Web Consortium (W3C) standard for data interchange that deals with the relationships between data, i.e., with the semantics embedded in it, and provides a foundation for publishing and linking these data. Within RDF graphs, entities are uniquely and globally identified using (http/s) Unique Resource Identifiers (URIs), enabling global interoperability. This enables frictionless interaction between your KG and the Linked Open Data cloud, which is a large and rapidly growing set of external information resources, some of which are published by government agencies.

Some aspects of building KGs can be automated using modern machine learning techniques, which can, given examples, automatically classify complex data (e.g., documents, images, videos) and define new classes without requiring explicit description. This classification establishes a new relationship between data points that belong to the same class — for example, documents that are real estate contracts — which then can be verified by domain experts and incorporated into a KG. The relations between the original data points are pre-calculated and the relationships themselves become an important part of the dataset, so each relation can be analyzed quickly and at scale. The choice of how to describe the relations and the ability to analyze fast and at scale is the key to new insights.

Marketing and Business Development Use Cases

Now that we have introduced the relevant technologies, let's look at some of the marketing and BD use cases.

USE CASE 1: OPPORTUNITY SPOTTING

PROBLEM

Law firms that monitor and analyze news and court dockets to spot business development (BD) opportunities need to more efficiently filter out noise and detect the few truly relevant events. Firms also need more efficient ways to integrate information from various internal and external sources into BD alerts sent to partners and BD managers that will enable these professionals to decide whether and how to pursue an opportunity.

SOLUTION

"Knowledge graphs are the next forefront in current awareness since they add the missing element from many current monitoring systems — context," says Leslie Lanphear, director of research and information services at Akin Gump.

Instead of using the same BD opportunity alert template that contains static data fields about a company for all alerts, firms can use ontologies and KGs to customize alerts by dynamically suggesting data fields and questions about a company or opportunity to include in the business development alert. "Systems that create an alliance between human and artificial intelligence will help our research services department shift to a more

proactive mode, and anticipate what questions about an opportunity a partner would want answered," says Scott Bailey, director or research services at Eversheds Sutherland.

USE CASE 2: AUTHORING BETTER CLIENT ALERTS

PROBLEM

When writing client alerts to keep clients well-informed about regulatory developments, authors of the alerts need to manually search multiple systems to find relevant information, such as past alerts, background reference information and profiles of relevant attorneys.

SOLUTION

With systems powered by KGs, authors of client alerts can leverage their firm's proprietary knowledge without manual searching. A system that uses KGs can surface relevant information from internal sources, such as the Document Management System (DMS) and internal publications. "A system powered KG will power automated matching to relevant content, enabling us to produce higher-value client alerts more efficiently and consider our readers' needs for intelligence," says Scott Leeb, the director of knowledge management at Fragomen, Del Rey, Bernsen & Loewy, LLP.

USE CASE 3: FOSTERING CROSS-SELLING AND STAFFING PROJECTS

PROBLEM

There are many instances in which law firms need the ability to match the right attorney to the right project or matters, such as to foster cross-selling opportunities or to identify the appropriate attorneys to include in an RFP. However, it can be difficult to identify the appropriate attorney when they need a specific set of attributes. For example, many law firms would find it difficult to identify which lawyers in the firm possess all of the following qualifications: speak Spanish, are veterans and have worked on pharma industry deals over $1 billion in the last three years. Limited and potentially outdated attorney profile information reduces cross-selling and other opportunities. Lawyers are more likely to reach out to unfamiliar colleagues if attorney profiles display rich information about their interests, relationships, and experience with a given industry, client and type of matter. Unfortunately, most lawyers often will not spend the time to update their profiles manually.

SOLUTION

Systems that leverage a combination of KGs and Natural Language Processing (NLP) techniques can automatically build more accessible and useful attorney profiles — which are easier to match to project and matter descriptions — by enriching them with standardized metadata. KGs can help systems extract structured information

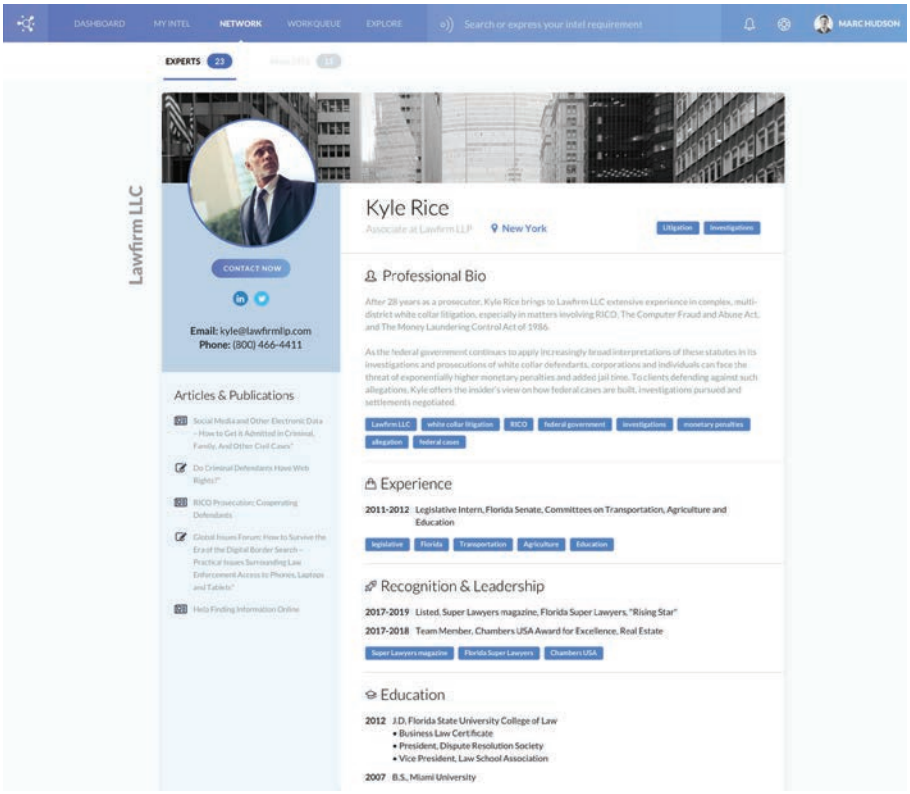

A KG-powered system could add skills (in blue) to attorney profiles automatically using NLP.

from attorney profile summaries and work product documents written in plain text. A KG-powered system can help suggest skill additions to attorney profiles, e.g., "attorneys with this skill have a 15% higher chance of getting a new matter of this type." Artificial Intelligence (AI) is critical because manually tagging documents and adding metadata to lawyer profiles would be an extremely time consuming and likely cost-prohibitive task.

CONCLUSION

Law firms often lack a holistic view of their clients and prospects. They cannot run sophisticated searches because their firms haven't created a repository of integrated data from Customer Relationship Management (CRM) software and other sources of data (e.g., finance; DMS; external sources of company, news and litigation data; internal research reports; thought leadership publications; CLE events) to obtain results with relevant content grouped together. KGs can help firms generate actionable insights to improve conversion, retention and advocacy rates. "With a more holistic picture of a client or prospect, we can discover opportunities and risks pertaining to client relationships that might otherwise go unnoticed," says Kelly Enache, CMO at Saul Ewing.

While KGs may sound like a futuristic pipe dream to some, they are a proven set of technologies used by many organizations outside of legal. A few law firms are implementing their first applications powered by KGs. They realize that the volume and variety of information they need to integrate is growing dramatically. Law firms need to know about the data's structure; a KG and semantic standards to describe the structure of information in the graph can support reasoning and inference. An ontology that defines the meaning of the data makes queries, algorithms and analytics more powerful.

Once law firms build KGs for initial use cases, they will be able to leverage and extend them for other purposes and feed them with automatically extracted information from a variety of content sources and systems. Law firms that take advantage of this technology will supercharge their ability to go from data to information, information to knowledge and knowledge to strategically-valuable insights. They will be able to distill the most important expertise from a sea of data and deliver exactly what the client needs, when they need it. ■

David Kamien *is the CEO of Mind-Alliance Systems (www.mind-alliance.com) and has spent 20 years advising law firms, legal departments and corporations in the fields of knowledge management, business development and competitive/market intelligence. David is the product owner of the AI-powered MindPeer business development intelligence platform. David has a law degree from the College of Law and Business and a B.A. from Hebrew University. Contact David at david@mind-alliance.com to sign up for his free weekly intel brief about legaltech innovation.*

The content in this feature correlates directly with the **Technology Management** domain in the LMA Body of Knowledge. To dive deeper into this subject area, head to these domain here: **http://bit.ly/LMABoKTechMgmt**.