RWS

Your guide to search experience optimization

Because it should be easy for your customers and employees to find what they need

Do you have an information findability problem?

If you were an **insurance provider**, what would you do if 90% of your employees were spending between 30 minutes and an hour, *every day*, searching for information? Apart from the obvious productivity implications, this means customers are hanging on, waiting for answers – and possibly choosing a competitor that provides a better experience.

And what if it's your customers who can't find what they need? Maybe you're a **provider of healthcare information to medical professionals** or you **sell electric cars** – whatever your business, you're not alone in wanting to do more to help customers find the products and services, support content and other information that they're looking for faster.

Why is findability so hard to crack?

Whether it's employees or customers searching for answers, the reason they too often struggle or fail is the same: organizations are doing a poor job of making their information findable. This is down to several issues that exacerbate one another:

Volume. The sheer amount of information organizations generate makes it harder to find the right information, whether by search or by navigating a site or system.

Structure. There's little or no strategy or control over how information is organized or categorized, which is critical for simplifying search and navigation.

Terminology. If information is categorized at all, there's no agreed common terminology for doing so (for example, similar content is classified as being about 'bikes' by some and 'bicycles' by others, without a system for recognizing that these are synonyms; which means that someone searching for 'bikes' won't find the information about 'bicycles').

Silos. Too often the information relevant to a particular search is stored in a variety of disconnected systems – with obvious implications for findability. And when the answer is elusive, figuring out who to contact for help can be just as hard. Often the result is that similar information is created in different places by different people – creating a new problem for others who discover these different versions and don't know why they exist or whether the differences matter.

Silos are a challenge when resolving a customer inquiry[1]

65%

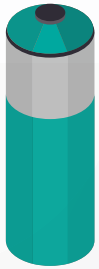

say silos between departments prevent effective internal communication

56%

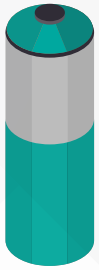

say there are too many different systems/tools to consult

1 Harvard Business Review Analytic Services, Improving the Customer Service Experience, September 2021

Customers expect more

Despite how common these issues are, we know that poor findability can be fixed. After all, organizations like Google and Amazon make it really easy for us to find what we need.

This raises expectations: why can't it always be this good? We should be able to find what we need, without effort. Certainly as customers – but why not also at work?

Employees need better

Since the pandemic struck, many more people are working from home more often than before. And when remote workers need to know or find something, they can't just ask those around them. They can't just wander over to the desk of the office 'guru' who's been there 20 years and knows everything. This makes it ever more important for them to be able to search for information and find it quickly and reliably – which includes being able to trust that the information found is accurate and up to date.

The gold standard of findability

Based on our best experiences with search interfaces, we know what 'excellent' looks like. It's when the systems that a searcher interacts with can almost read their mind:

- **Autosuggestion:** offering relevant autocomplete options as they type.
- **Personalization:** returning results and recommendations that show an understanding of the searcher and their situation, for example their location, age, role or recent online activity. And only presenting results that are relevant to the searcher.

What does it take?

To achieve this gold standard, you need specific capabilities in three areas:

- **Knowledge management**
- **Semantic AI**
- **Structured content management**

Any of these on its own will produce results and get you closer to the gold standard. But the real gains come when you bring them together, as we'll discover.

We'll also touch on how to make a good start:

- **Choosing the right technology**
- **What content to start with**

> **A customer's expectation of speed isn't shaped by the best in your competitor set; it's established by the best experiences they've ever had.**
>
> **Liz Barnsdale**
> Managing Director, Accenture Interactive[2]

[2] Harvard Business Review Analytic Services, Improving the Customer Service Experience, September 2021

Knowledge management best practices for great findability

Whether or not you're familiar with the practices and concepts of knowledge management, you'll know instinctively that the better organized information is, the easier it is to search or navigate. That's why most companies use at least some basic form of classification to help organize their information.

Beyond basic: taxonomy and knowledge graphs

Unfortunately, basic classification won't get you close to the gold standard for search experience optimization and findability. You need at least a formal taxonomy, and preferably an enterprise knowledge graph (the two aren't exclusive: taxonomies provide very valuable input to the development of knowledge graphs and make them more effective).

Definitions

Taxonomy is a scheme of classification, especially a hierarchical classification, in which things are organized into groups or types.

Knowledge graphs offer much greater flexibility than taxonomies in the relationships you can create to link objects and concepts, and in allowing metadata for taxonomy terms (metadata about metadata). This flexibility allows knowledge graphs to go a long way towards encoding the semantics, or meaning, of information through its attributes and relationships with other information.

Knowledge organization continuum

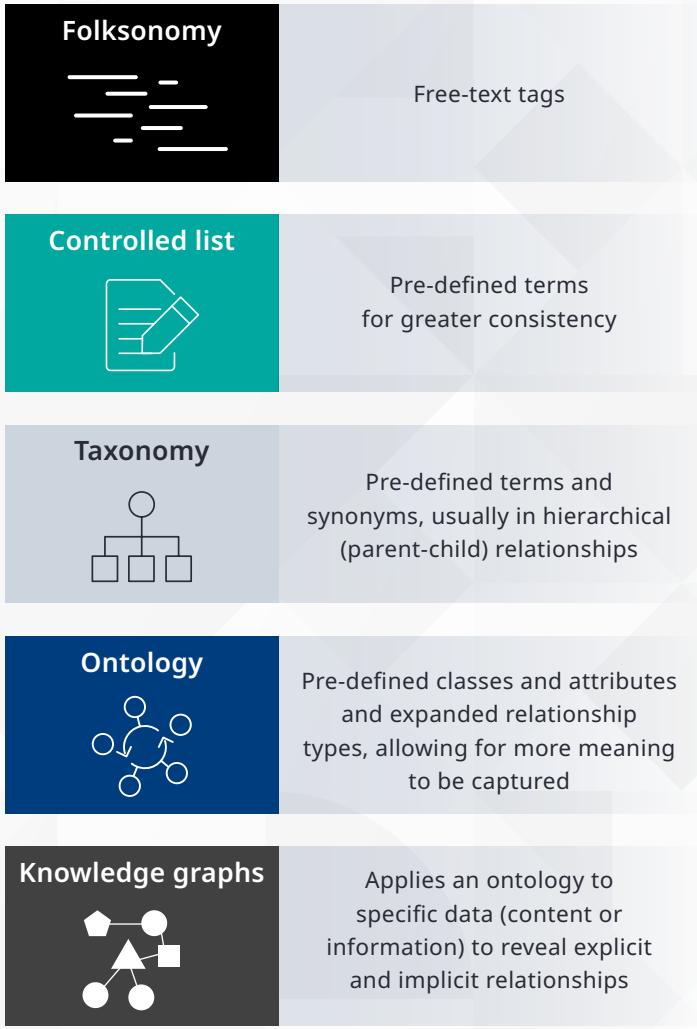

Why a knowledge graph?

Taxonomy is a great starting point for improving search. It formalizes the language you use to categorize content: allowing you, for example, to recognize that 'bike' is a synonym for 'bicycle' – and, indeed, for 'motorbike'. And it organizes information in a straightforward way. When a website lets you filter your searches to narrow in on what you want – a gravel or mountain bike rather than a commuter or cargo bike, for example – that's taxonomy at work.

But taxonomy's narrow scope for expressing relationships does limit its effectiveness. If someone in Copenhagen searches for 'bike repair', how can we know if they mean 'bicycle' or 'motorbike'? If all we have is taxonomy, we can't. But a knowledge graph can capture the fact that 'bike' is more likely to mean 'bicycle' than 'motorbike' in Denmark, greatly improving our chances of returning relevant results. We can also return the names of bike repair shops in Copenhagen rather than in, say, Denver, because our knowledge graph also captures the relationship of 'is located in' (and 'has a mechanic') for bike shops.

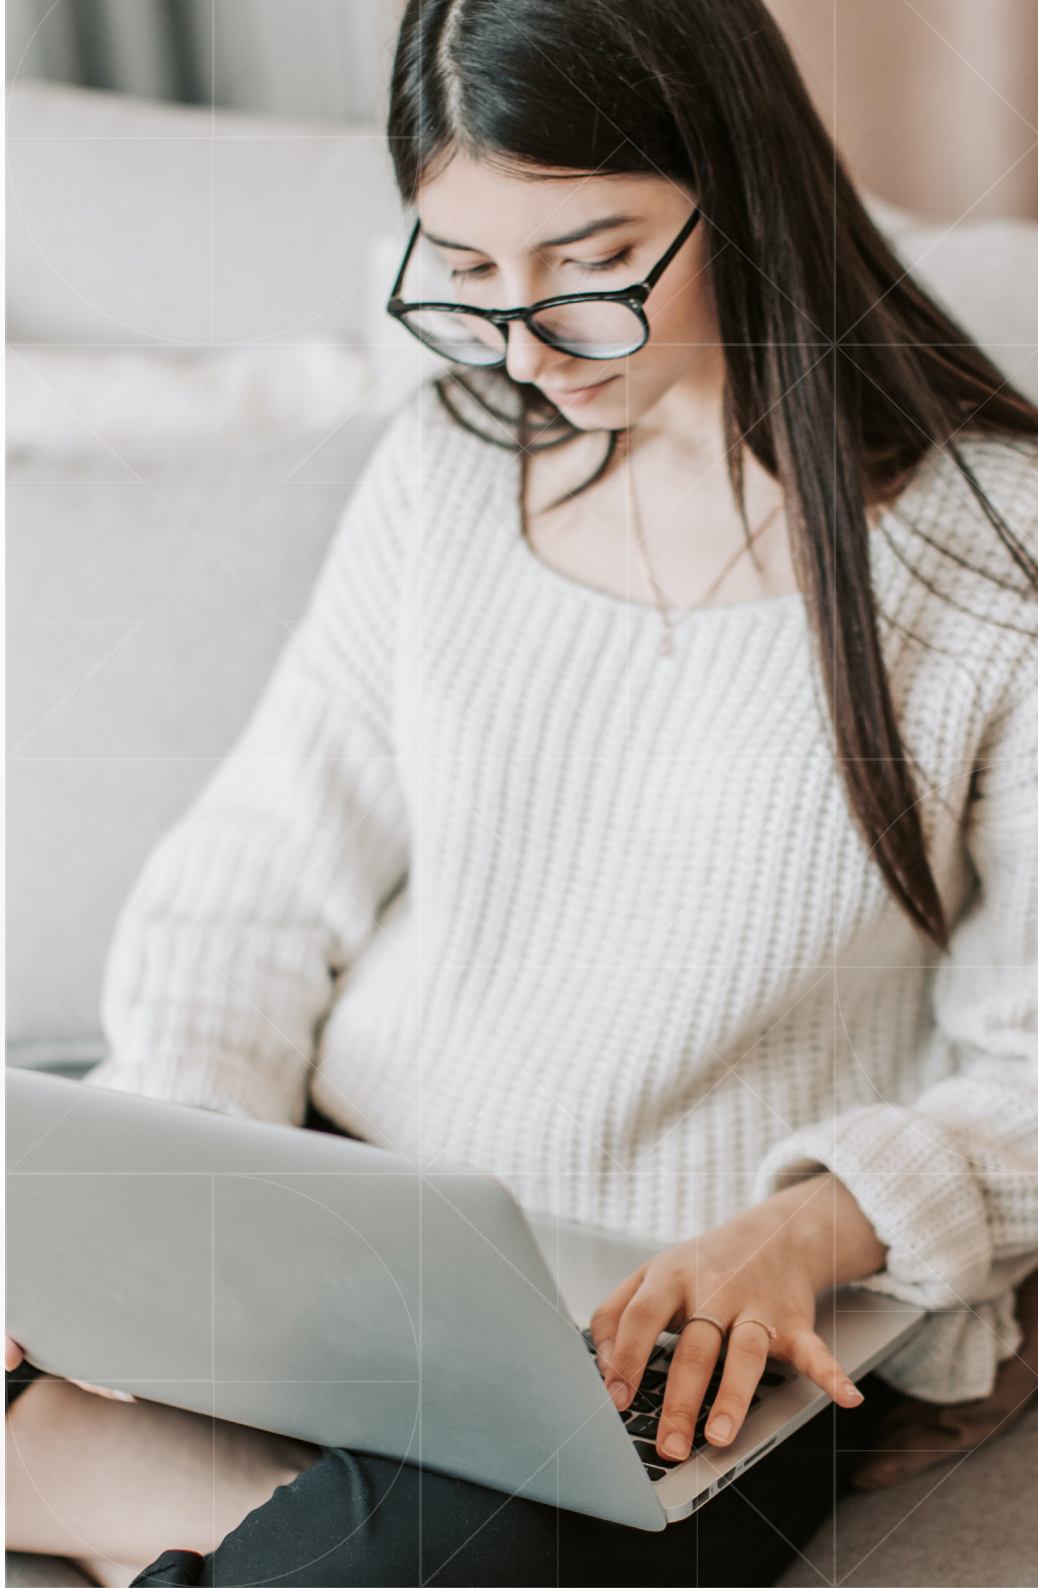

Breaking down barriers to learning

Companies can also use the power of knowledge graphs to connect different teams and break down internal communication barriers. Simply by capturing the owners of specific content or topics in your business, you can help employees identify who to contact if they can't find what they need.

There are many, many more ways in which the flexibility of knowledge graphs to capture all sorts of relationships and semantic information improves our ability to hone in on the right information. And as a knowledge graph grows – or we connect different knowledge graphs – connections that may have remained hidden can become exposed, allowing for new inferences.

The example here is simple but illustrates how we can infer that you need to go to Paris to see the Mona Lisa, without ever explicitly locating the Mona Lisa in Paris. When we add AI to the mix, as we'll discover, this forms the basis of more relevant, personalized search results and recommendations.

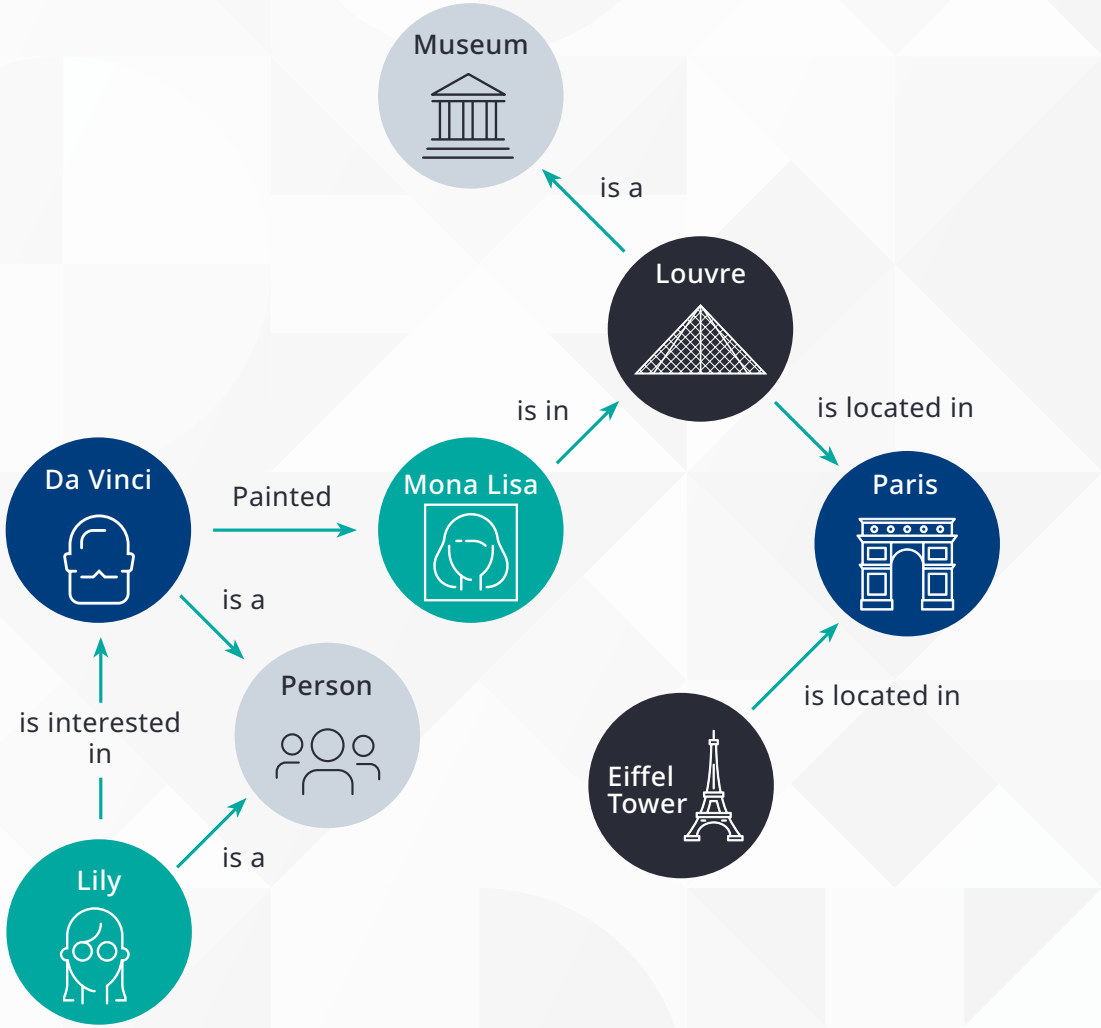

Semantic AI: the future of findability and online personalization

Not too long ago, the only way to personalize content for online visitors was to hard-code different scenarios ('for the same search, respond with X if the searcher is in America and she's 35-50 years old, respond with Y if she's 18-35, Z if she's a he …'). Obviously this approach requires an awful lot of work for a very limited capacity to respond to personal details about people.

Artificial intelligence (AI) has completely changed the game.

The strengths and weaknesses of machine learning

Using machine learning (ML), a search solution can discover what is helpful to people with certain characteristics and behaviours, and use this knowledge to improve results for others with similar characteristics and behaviours.[3] It can autocomplete searches, personalize results and offer relevant recommendations – all of which characterize the gold standard for search experience optimization.

Job done. Right?

Well, not quite. The power of machine learning to be relevant can be limited by one or more of these issues:

- Lack of enough past behaviour to uncover user intent, including inability to 'understand' ambiguous or conversational language
- Limited known characteristics of a searcher – if, for example, they refuse cookies
- Content not organized and categorized well enough for reliable identification of specific information.

[3] This works better if you have good feedback loops that encourage searchers to rate the response they get and say why it was good or bad.

Stepping up to semantic AI

With semantic AI, you can go a long way to overcoming these limitations. While past behaviour and information about the searcher will always improve results, their absence is no longer such a deal-breaker if you add the power of semantic AI to a solution that can already learn and adapt.

Semantic AI is AI that appears to understand what content means. This 'understanding' comes from a combination of corpus analysis (a trained AI scanning content to identify what it's about) and the relationships and attributes that are captured by a knowledge graph (see the definition of knowledge graph on page 6). In simple terms: give the right kind of AI-powered solution access to a knowledge graph, and you get semantic AI. The more comprehensive and granular the knowledge graph, the more the AI will 'understand' and the better it will perform.

Semantic AI is the most reliable way to achieve the gold standard for search experience optimization and findability. It's the secret to providing personalized results and recommendations as well as dynamic, relevant autosuggestions and concept search results to your customers and employees[4].

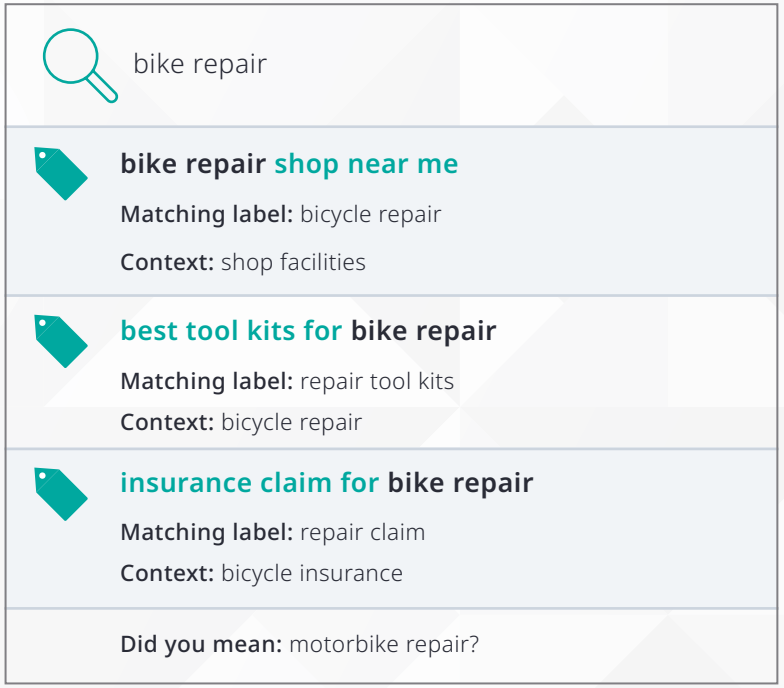

[4] A concept search is one that retrieves information that is conceptually relevant to the search query, without requiring an exact language match.

Semantic AI makes the implicit explicit

Earlier we noted that knowledge graphs can expose connections not explicitly made, allowing for new inferences. And as we know, AI can be exceptionally good – often better than humans – at spotting patterns and connections that aren't obvious. This ability adds to the power of semantic AI to find relevant information and recommendations for searchers.

Another way that semantic AI can help with findability is by improving content categorization as content is authored or updated – using corpus analysis and its access to the knowledge-graph to suggest tags that the author might not have thought of, enabling synonyms and related concepts to be captured. This also helps to enforce the consistent use of agreed taxonomic terms.

Semantic AI isn't just for search

While we've been focusing on search, semantic AI is just as valuable for the related 'self-service' tools of chatbots and voice assistants, enabling you to deliver personalized, human-like interactions without the involvement of a human.

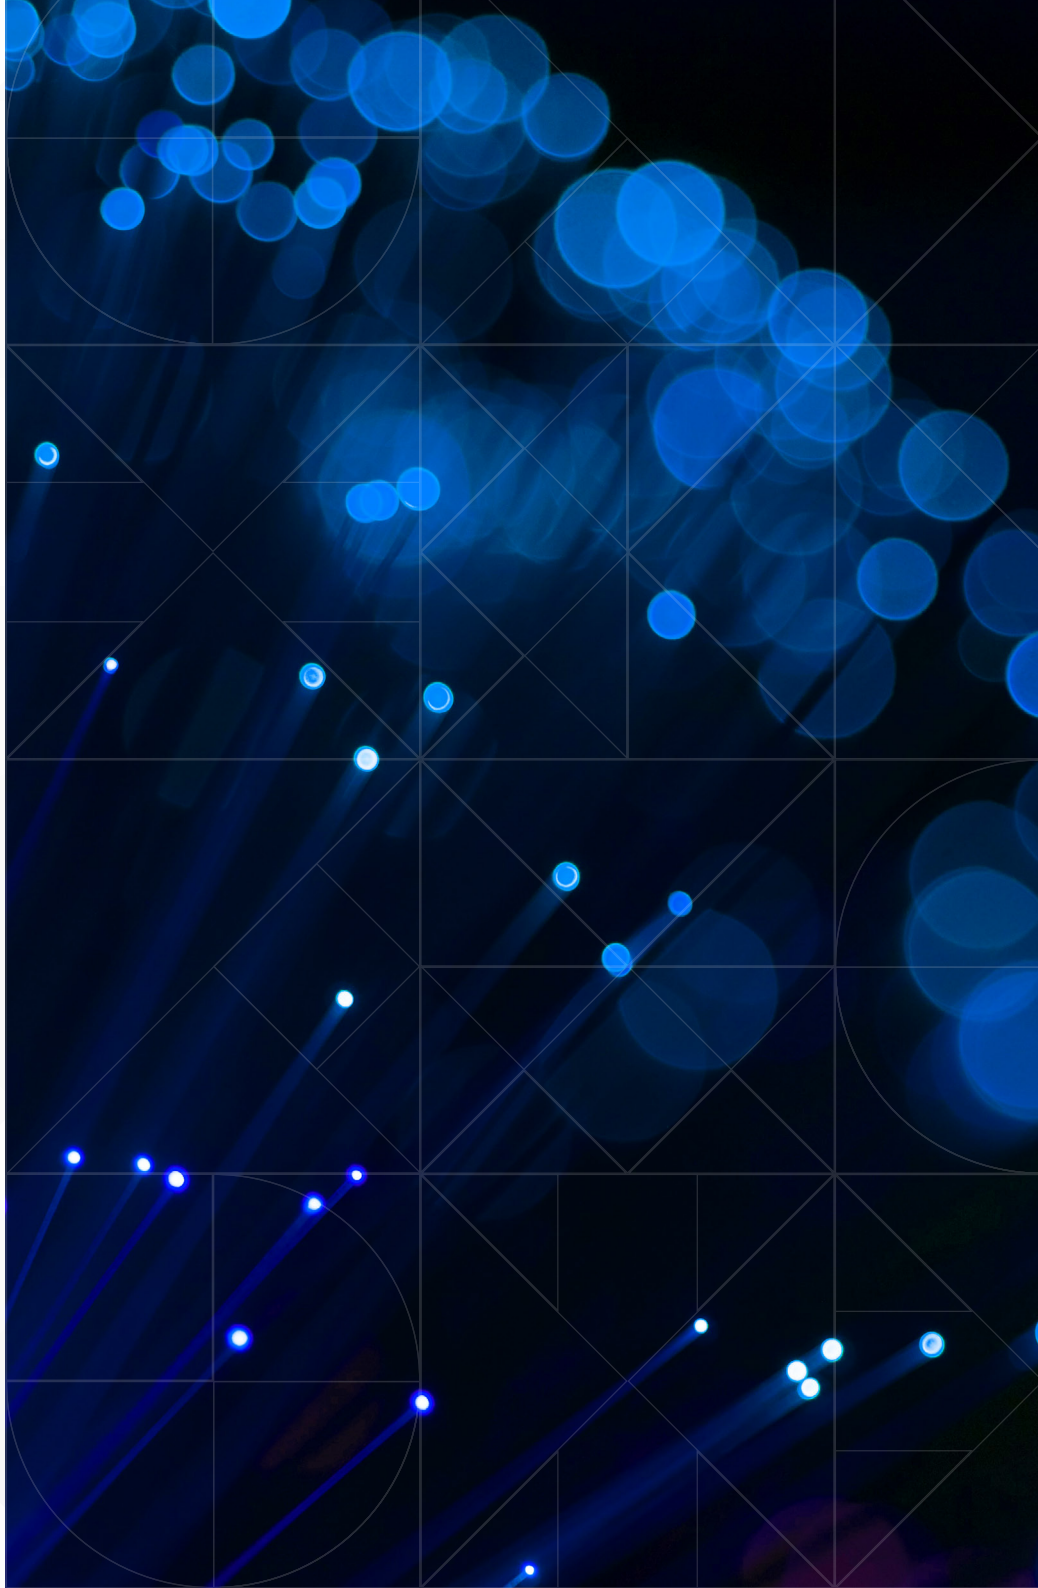

The value of structuring content through component management

Imagine you sell hundreds of motorbikes and each has a manual that you publish online (as PDF documents or html web pages). You can categorize the manuals by the bike model they relate to and, depending on the sophistication of your classification system, by the type of content they contain, or the types of queries they might be relevant to.

Specificity = greater relevance

Now what happens when a searcher comes to your site and searches for 'engine making weird noise'? After you've had them provide their bike model, you'll be able to direct them to the manual – or perhaps directly to the section dealing with engine troubleshooting. That's good; but not nearly as good as being able to ask them questions or provide filters to narrow down the type of noise – is it a gurgling, ticking, clattering, knocking or popping sound – and directing them to the content that is more specifically relevant.

This ability to get more precise – and do so automatically and reliably when responding to searches or structuring content for navigation – is determined by how granular your content categorization is. And to get more granular with content categorization, you need to be able to break your content into smaller components.

Enter the component content management system (CCMS)...

A CCMS is designed to hold content as individual components rather than as whole documents. A component could be a word, phrase, paragraph, series of paragraphs, image, video, table, or any other 'piece' or 'module' of content. The CCMS is all about:

- Formalizing and standardizing the structure of content as a set of related components
- Simplifying the management of modular content, including categorizing (or tagging) components, updating them (managing different versions in time), translating them (managing different language versions), and pulling them together as necessary into traditional documents or other formats

This approach really comes into its own when dealing with content that needs to be reused a lot – product manuals, for example, will have many similarities across different models in a product family. Much of the content you want to make more findable by customers and employees will fall into this category, including product information, knowledge bases for customer support, policies and legal documentation.

There are at least three reasons why use of a CCMS will help you achieve the gold standard for search experience optimization and findability:

1. More powerful knowledge management

As already mentioned, a CCMS lets you get more granular with categorization, resulting in more powerful knowledge graphs and, ultimately, more relevant results.

2. 'Single source of truth'

When an insurance company needed to update the down-payment amount in its policies, it took six people searching for a week to find and replace 26 instances across its 35,000-document library – only to later discover more instances that were synonymous derivatives[5].

Now put yourself in the shoes of an employee at this company trying to answer a customer question about a down payment. After some digging, you find what you think is the right policy document. But can you be sure? There are different versions and you can't tell if they're duplicates or newer or older. And even if there's one clear version, how confident are you that the information in it is accurate?

By allowing you to manage content once and giving you full, flexible control over version management, a CCMS eliminates this kind of problem. It gives you a single source of truth, the 'master' content that you only ever have to update once, before pushing it to wherever you're (re)using it. It destroys content silos and unwitting content duplication. Not only can people and systems find the information they need, they no longer have to worry that the information may be out of date.

At least once a day, 16.5% of workers will create a new information asset – only to learn that a similar asset already exists.[5]

3. Headless CMS: deliver to any format or channel

The right kind of CCMS can also be a headless content management system, which means it separates content from its presentation or formatting. This makes it much easier to push content to different channels and choose different formats for each channel – which simplifies the job of providing information to customers and employees wherever and however they want to consume it.

[5] IDC, **The Future of Knowledge Management: Agile, Governed, and AI Ready Componentized Content Services**, Aug 2021

Choosing the right technology

To achieve the gold standard for search and findability you've got to invest in technology that helps you create and maintain a knowledge graph, manage your content as components, and apply semantic AI to search and other applications.

It is possible to patch together a knowledge management (KM) system, CCMS and AI-enabled search platform to create the solution you need. But implementing, maintaining and integrating different systems is costly and time-consuming for any business.

Intelligent content management

It's far more efficient to choose a single solution that gives you the core CCMS, KM and semantic AI search features that you need – all in one place, already working seamlessly together.

There's no standard way to refer to such a solution. We call it **intelligent content management** – not just because it uses semantic AI, but because it helps you work smarter to get the right content to the right people at the right time in the right way.

What content to start with

No organization can apply the gold standard of findability to all of their content and data all at once. You've got to start somewhere, which means choosing what to prioritize.

Use your business objectives to guide you in this, especially if your aim is to improve findability for employees. For customers, focus on what drives them to contact your support desk or on what they search for on your corporate website, ecommerce site, or to get to you from search engines.

While it does take time to apply knowledge management best practices to content or to componentize it, it doesn't have to be as time-consuming or complex as you might think. This is especially so if you're using the right solution and don't have to wait for complex IT integrations in order to get started. We know businesses that have started delivering results within weeks with our intelligent content management platform.

Leading the way

RWS is **recognized by Ars Logica** as the provider of the number-one CCMS, and **also by IDC** as a leader in content management systems for authenticated digital workspaces.

Beyond the tech, we can also advise and guide you on knowledge graph best practices and help you get started with the right content.

Want to find out more?

To further explore the role and value of intelligent content management, read the IDC Technology Spotlight: **The Future of Knowledge Management: Agile, Governed, and AI-Ready Componentized Content Services.**

To discuss your specific situation and how we can help

contact us

About RWS

RWS Holdings plc is the world's leading provider of technology-enabled language, content management and intellectual property services. We help our customers to connect with and bring new ideas to people globally by communicating business critical content at scale and enabling the protection and realization of their innovations.

Our vision is to help organizations interact effectively with people anywhere in the world by solving their language, content and market access challenges through our collective global intelligence, deep expertise and smart technology.

Customers include 90 of the globe's top 100 brands, the top 10 pharmaceutical companies and approximately half of the top 20 patent filers worldwide. Our client base spans Europe, Asia Pacific, and North and South America across the technology, pharmaceutical, medical, legal, financial services, chemical, automotive, government and telecommunications sectors, which we serve from offices across five continents.

Founded in 1958, RWS is headquartered in the UK and publicly listed on AIM, the London Stock Exchange regulated market (RWS.L).

For further information, please visit: **www.rws.com**

© All Rights Reserved. Information contained herein is deemed confidential and the proprietary information of RWS Group*.
*RWS Group shall mean RWS Holdings PLC for and on behalf of its affiliates and subsidiaries.

tridion-launch-asset-self-service-ebook-rws-en-a4-181121