From microcontent to neurons

A practical guide for building a cognitive AI content supply chain for highly personalized user assistance

Michael Iantosca
Senior Director of Content Platforms
Michael.Iantosca@avalara.com

Avalara

Download full paper:
ThinkingDocumentation.com

About me

- Enterprise content systems architect
- 40th year in intelligent (structured) content
- Mentored by the inventor of structured markup language
- Cross-trained at IBM as an information developer and software engineer
- Designed and built multiple generations of intelligent content supply chains
- Multiple invention disclosures and patents in the content ML/AI space
- AI experience dating back to the early 90s with Expert Systems and LISP
- Formed the team at IBM that invented DITA in the late 90s

Artificial Intelligence and machine learning is simple and obvious.

Everyone gets it...

$$\mathbf{h}^\top \mathbf{M}_r \mathbf{t} + \mathbf{h}^\top \mathbf{r} + \mathbf{t}^\top \mathbf{r} + \mathbf{h}^\top \mathbf{D} \mathbf{t}$$
$$\mathbf{h}^\top \mathbf{M}_r \mathbf{t}$$
$$\sigma\left(\tanh\left(\mathbf{c}_r \circ \mathbf{h} + \mathbf{c}_r \circ \mathbf{h} \circ \mathbf{r} + \mathbf{b}\right) \mathbf{t}^\top\right)$$
$$g_{\text{left}}(\mathbf{h}, \mathbf{r})^\top g_{\text{right}}(\mathbf{r}, \mathbf{t})$$
$$\mathbf{h}^\top \operatorname{diag}(\mathbf{M}_r) \mathbf{t}$$
$$\mathbf{r}^\top (h \star t)$$
$$\sum_{j=0}^{l} p\left(\mathbf{h}, \boldsymbol{r}; \boldsymbol{c}_j\right) \cdot \boldsymbol{t}$$
$$-\left\|\mathbf{M}_r^1 \mathbf{h} - \mathbf{M}_r^2 \mathbf{t}\right\|_1$$
$$\frac{1}{2}\left(\mathbf{h} \circ \mathbf{r}\mathbf{t} + \mathbf{t} \circ \mathbf{r}'\mathbf{t}\right)$$
$$\mathbf{h}^\top \mathbf{M}_r \mathbf{t}$$
$$\mathbf{h}^\top \sum_{i=1}^{d} \boldsymbol{\alpha}_i^r \mathbf{u}_i \mathbf{v}_i^\top \mathbf{t}$$
$$\mathcal{W} \times_1 \mathbf{h} \times_2 \mathbf{r} \times_3 \mathbf{t}$$
$$\left(\mathbf{S}^k \operatorname{diag}\left(\mathbf{U}^T \mathbf{h}\right) \mathbf{V}^T \mathbf{r}\right)^T \mathbf{t}$$

Don’t you?

Agenda

- The challenges
- The goals
- The roadmap

Deliver the right content, to the right person, at the right time, and in the right experience.

We need to advance from providing reactive, *failure-mode* content to hyper-personalized, *pro-active* and assistive content

…and from more than just task-oriented, to scenario-oriented user assistance

"As an industry, we've mastered creating content for everyone, but for no one in particular."

mji

The challenge

- Lots of talk about AI/ML-enabled content, but a paucity of information about how to get there
- The focus has been on making content intelligent, but most discussions and efforts stop at reuse
- Too much focus on creating content intelligence, virtually none on *applying* that intelligence
- Too much focus on AI/ML bots, not on holistic cognitive content applications
- Virtually no discussion about the last mile: How to enabled cognitive content *retrieval, organization, and delivery* – treated as if it will somehow magically materialize

We can enable all these goals with a single intelligent content architecture and supply chain across the enterprise

Horizon I

- Fast and efficient content production
- Content reuse and repurposing
- Write-once, reuse-many single-source
- Efficient and cost-effective translation
- Improved organic and site search
- Omnichannel content delivery
- Improved content consistency

Horizon II

- Taxonomy autoclassification
- Self-service chat bots
- Seamless, integrated enterprise content experience (silo busting)
- Increased revenue through conversion rate optimization
- Personalized content
- Cross-sell/up-sell
- Robotic content generation

Horizon III

- Dynamic one-on-one personalization
- Scalable chat bots with *precise* answers and dynamic recommendations
- Pro-active user assistance
- Hyper-personalized content for multi-task scenarios
- AI-driven content discovery and reuse during creation (cross-silo) / ML-assisted reuse
- Autonomic user assistance - self-healing and adaptive content

Intelligent content → AI (non-ML) content → Cognitive ML content

Unified enterprise content experience and false prophets

Much has been tried - and failed over the decades:

- Templates
- Paragraph style standardization
- The “golden DTD”
- One CMS to rule them all!
- The canonical web portal
- Mashup portals

The semantic content maturity model (SCMM)

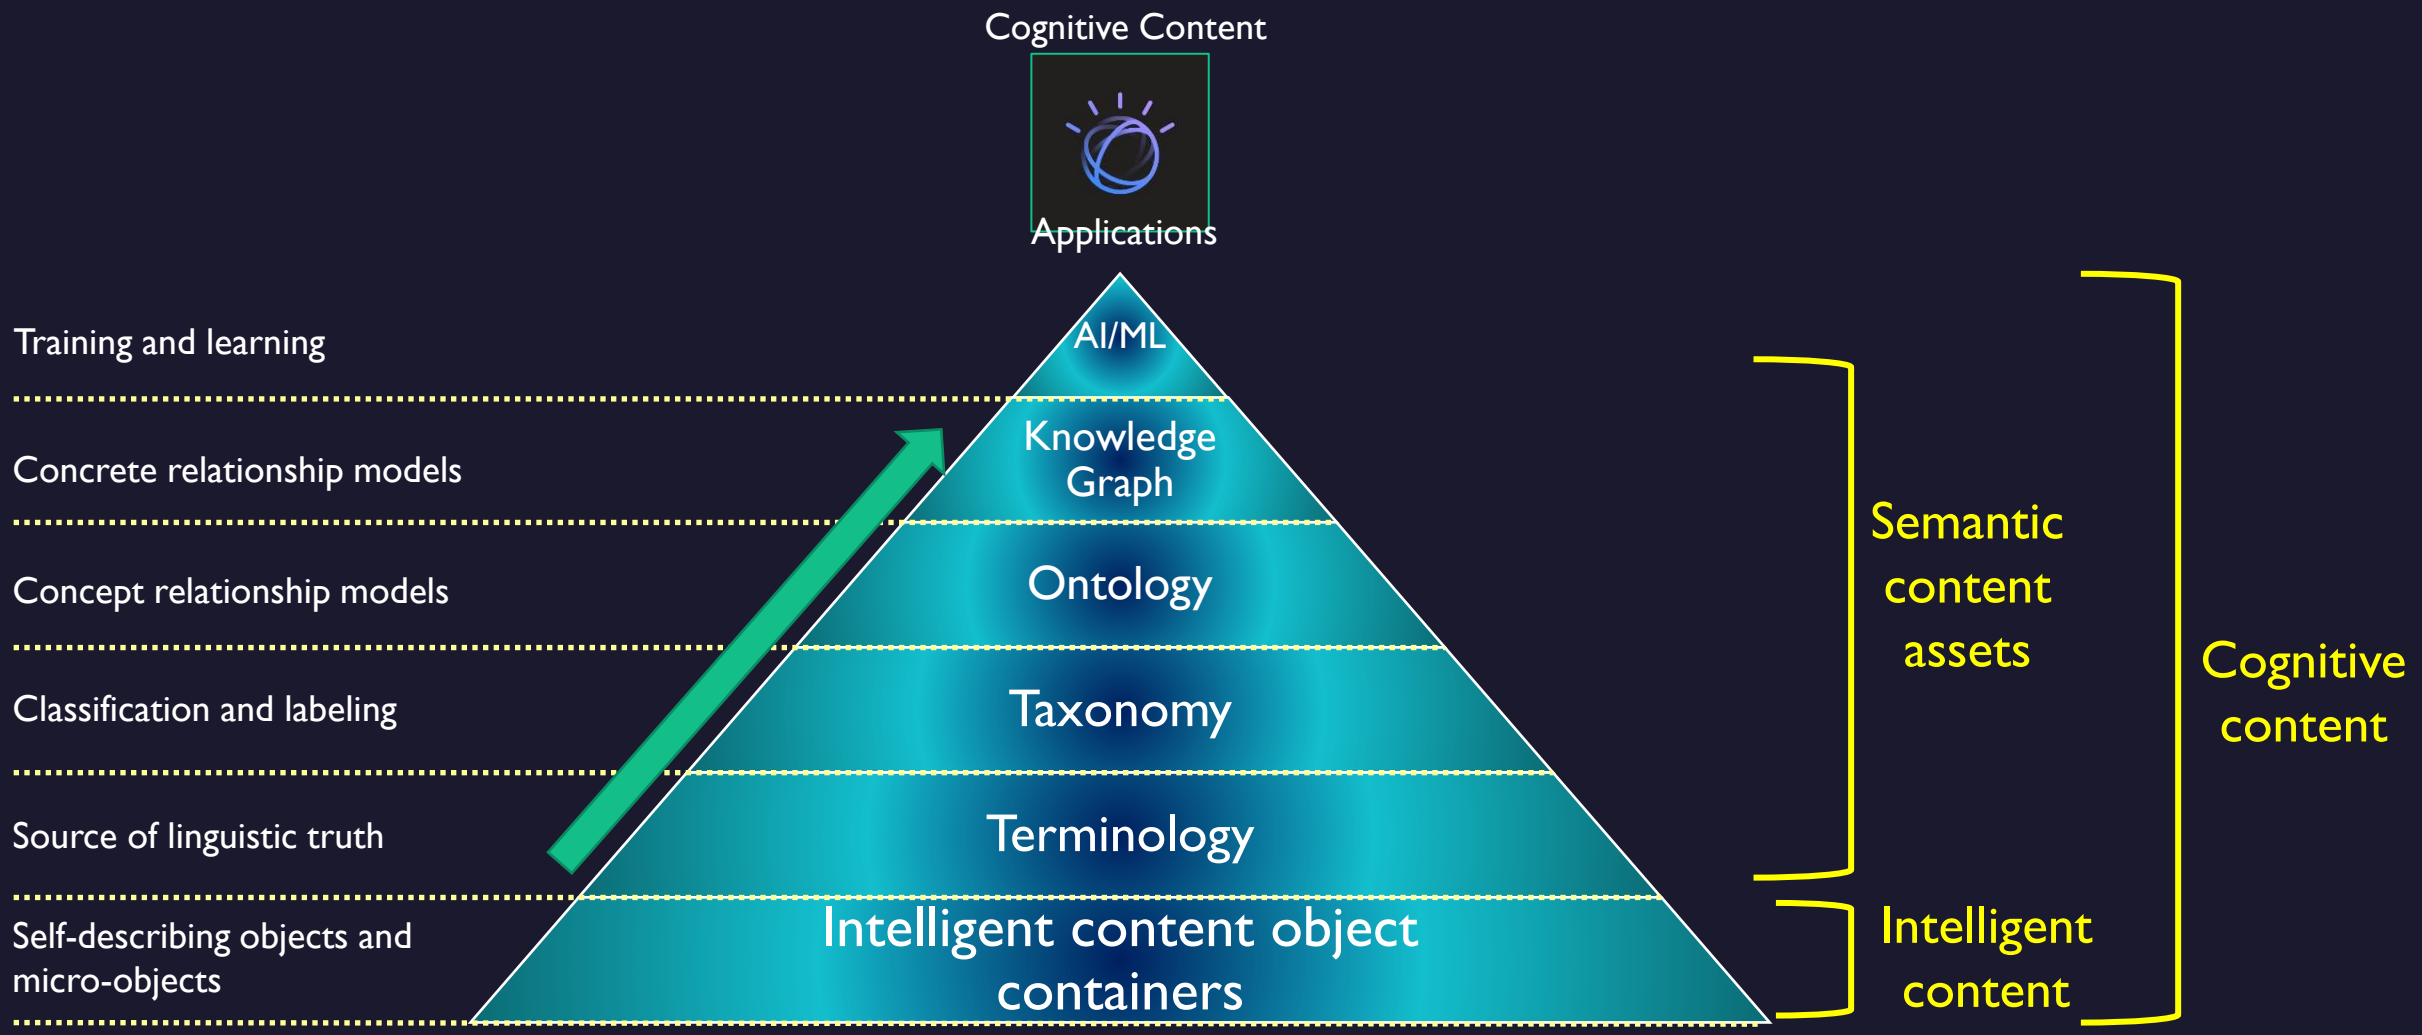

The semantic content maturity model (SCMM)

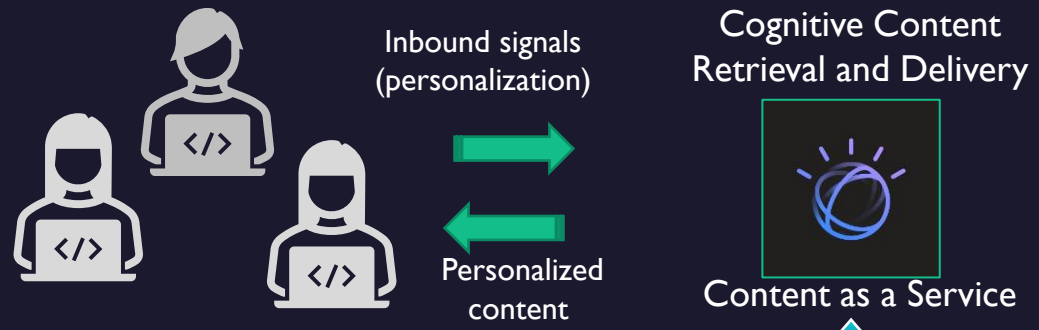

Inbound signals (personalization)

Personalized content

Cognitive Content Retrieval and Delivery

Content as a Service

Training and learning

Concrete relationship models

Concept relationship models

Classification and labeling

Source of linguistic truth

Self-describing objects and micro-objects

AI/ML

Knowledge Graph

Ontology

Taxonomy

Terminology

Intelligent content object containers

Semantic content assets

Intelligent content

Cognitive content

What is a cognitive content supply chain?

A cognitive content supply chain is a strategy, an architecture, and an operational model that enables dynamic, machine-based retrieval, assembly, and delivery of non-linear content objects to provide humans or machines with knowledge that is based on predictive relationships between content objects and inbound signals.

mji

Content object containers

- We need to advance beyond the component content model – to object-oriented
- Ideally a document object model such as DITA or JSON, but not exclusively
- Ideally inheritance-based (DITA) for maximum intelligence and flexibility
- Semantically descriptive wrappers
- Preferably typed content objects
- Portable metadata – intelligence travels with the content objects
- Microcontent / microdocuments

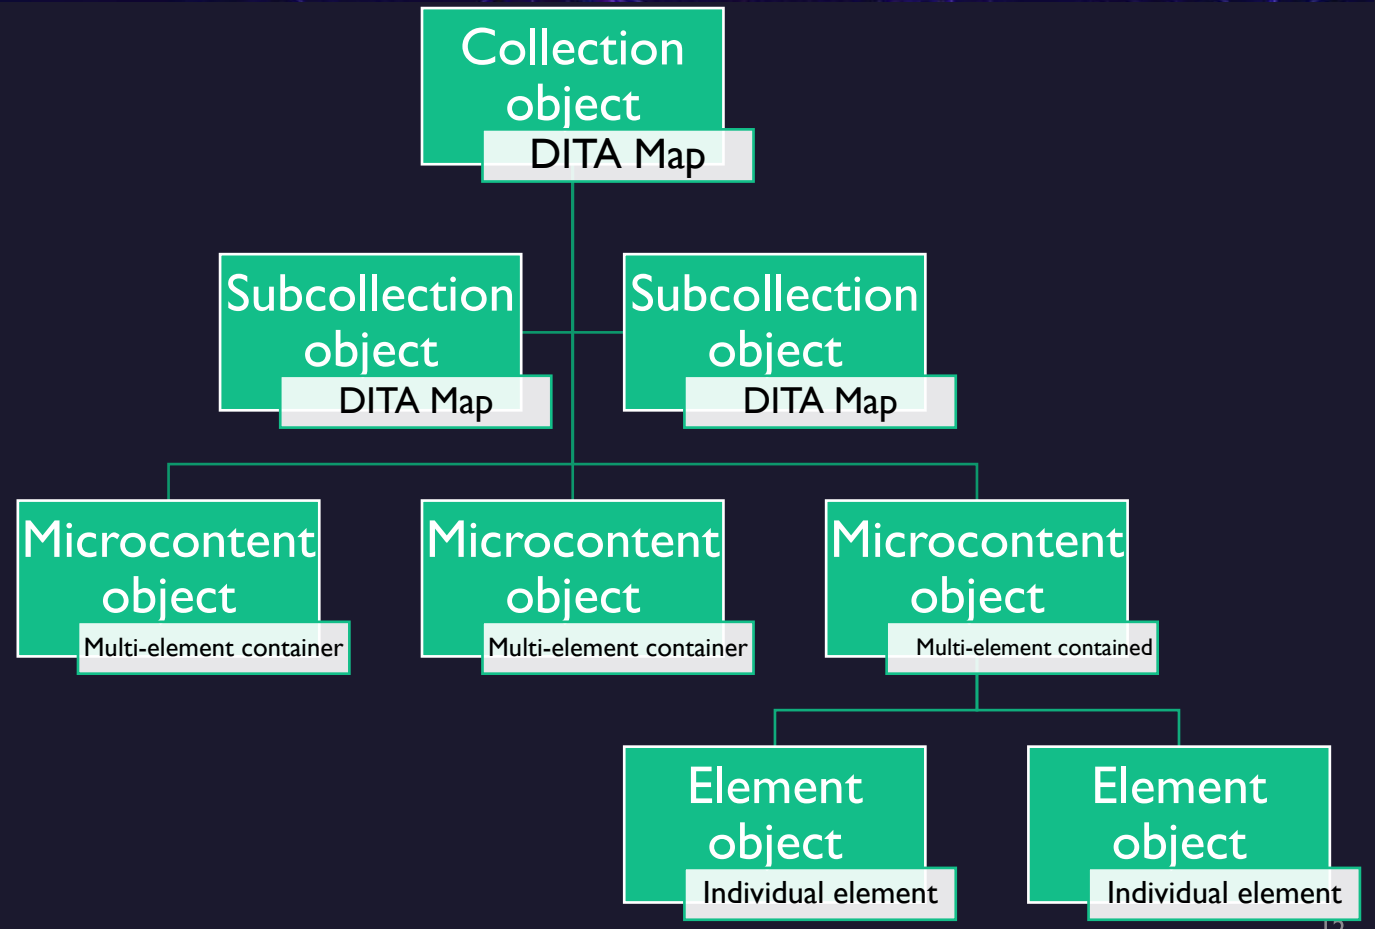

Enterprise harmonization of content types

- Collect all the names of content types from across the enterprise
- Include all content domains - from pre-sales through post-sales
- Define and manage it as one of your key enterprise content taxonomies ← This becomes an important semantic asset we can use later.

- Define each content type in detail
- Assign content type classification labels to your content objects
- A DITA concept, task, reference and others are not content types, they are *information* types

Define and map content journey and experience

Journey Stages	Discover	Learn	Try	Buy	Use	Advocate
Experience	Explore		Evaluate	Embrace	Extend	Expand

- Define the *macro* customer content journey
- Becomes another content taxonomy
- Apply it as *temporal* intelligence to content objects via taxonomy labels
- It enables the "…*at the right time*" element of our content strategy
- Map all your content types against the phases – include pre-sales through post-sales content types
- Do it for every product or service offering
- It is also useful for planning content and identifying overlaps and gaps across the enterprise
- *Workflow patterns* and *micro-workflows* (micro-journeys) are also needed. Workflow patterns are multi-task sequences for a specific goal.

Terminology governance

words matter, and they *really* matter with NLP and NLU

- Active terminology management
- Foundational, not optional
- Critical to align and harmonize content with:
  - Taxonomies
  - Ontologies
  - Knowledge Graphs
- Consistent terminology improves organic SEO and accelerates taxonomy autoclassification
- Use with computational linguistic services (term mining and in-editor author assist)

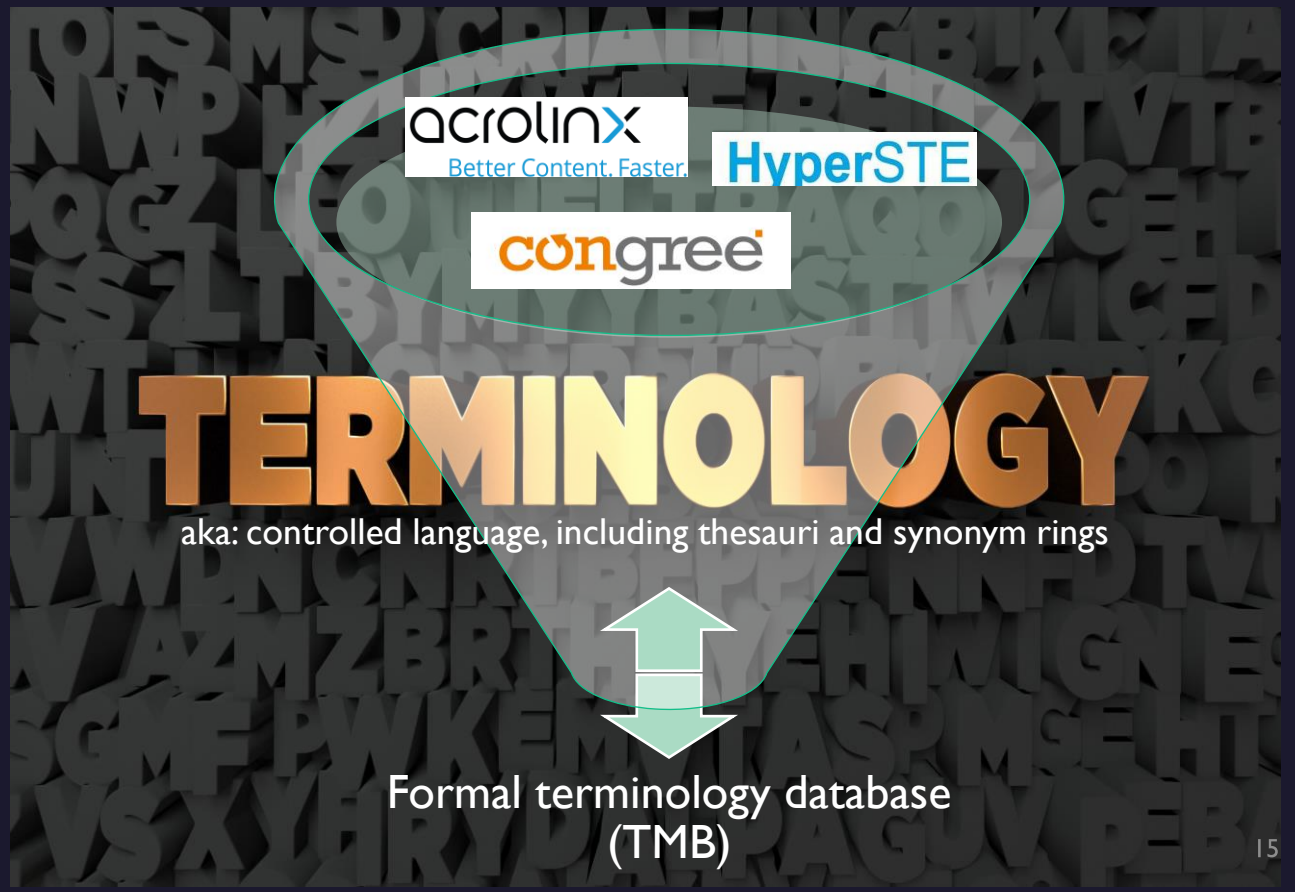

Computational linguistics services

Active content governance

- Content quality and consistency isn't optional
- Real-time, assistive AI editorial authoring assist
  - Accurate and consistent product and feature names
  - Consistent use of business terminology
  - Consistent tone of voice
  - Simplified and consumable
  - Accurate grammar and consistent writing style
- It's also essential to harmonize content before translation memories are generated to minimize translation and re-translation costs
- Helpful for non-professional content contributors

Ensure Avalara "speaks" with one consistent voice

Findable	Accurate	Measurable
Clear	Engaging	Effective
Consistent	Credible	Efficient

Build and manage enterprise content taxonomies

- Taxonomy is a classification of things
- If you've navigated eBay, you've used one
- Is a grouped or hierarchical controlled vocabulary
- Apply taxonomy labels to our object containers
- A few key taxonomies for AI discovery and retrieval:

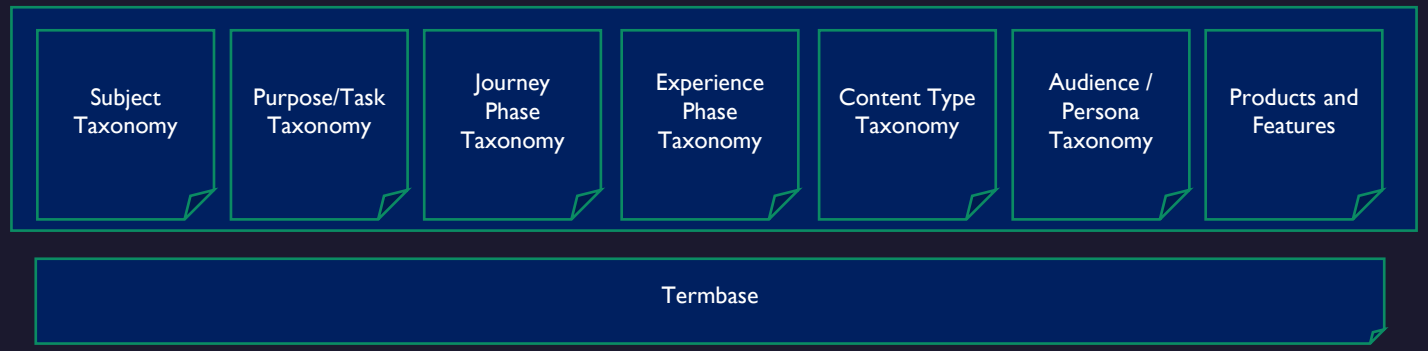

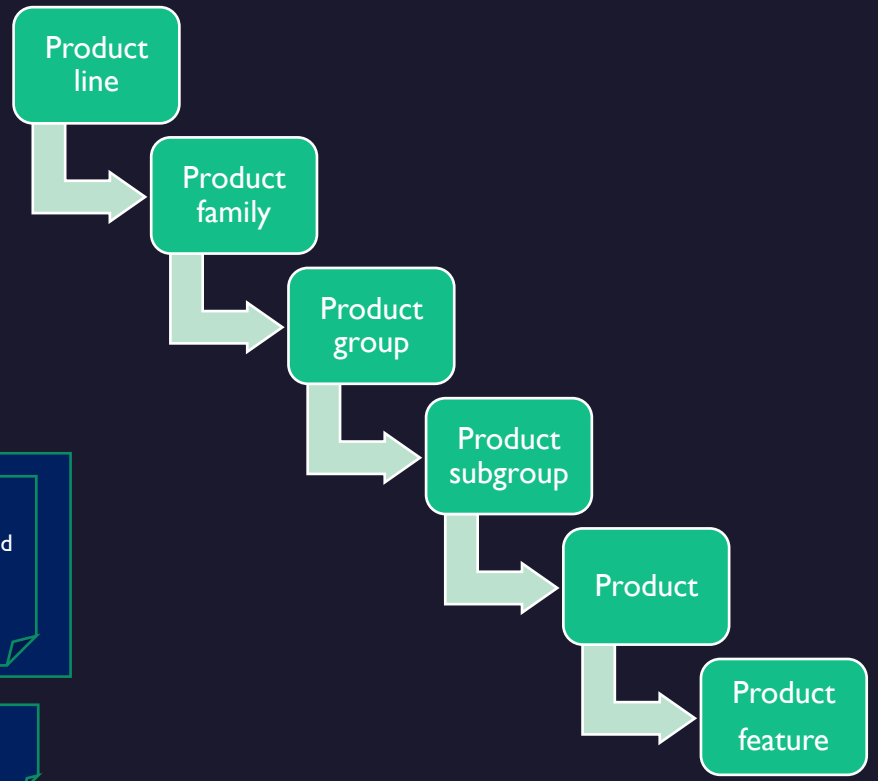

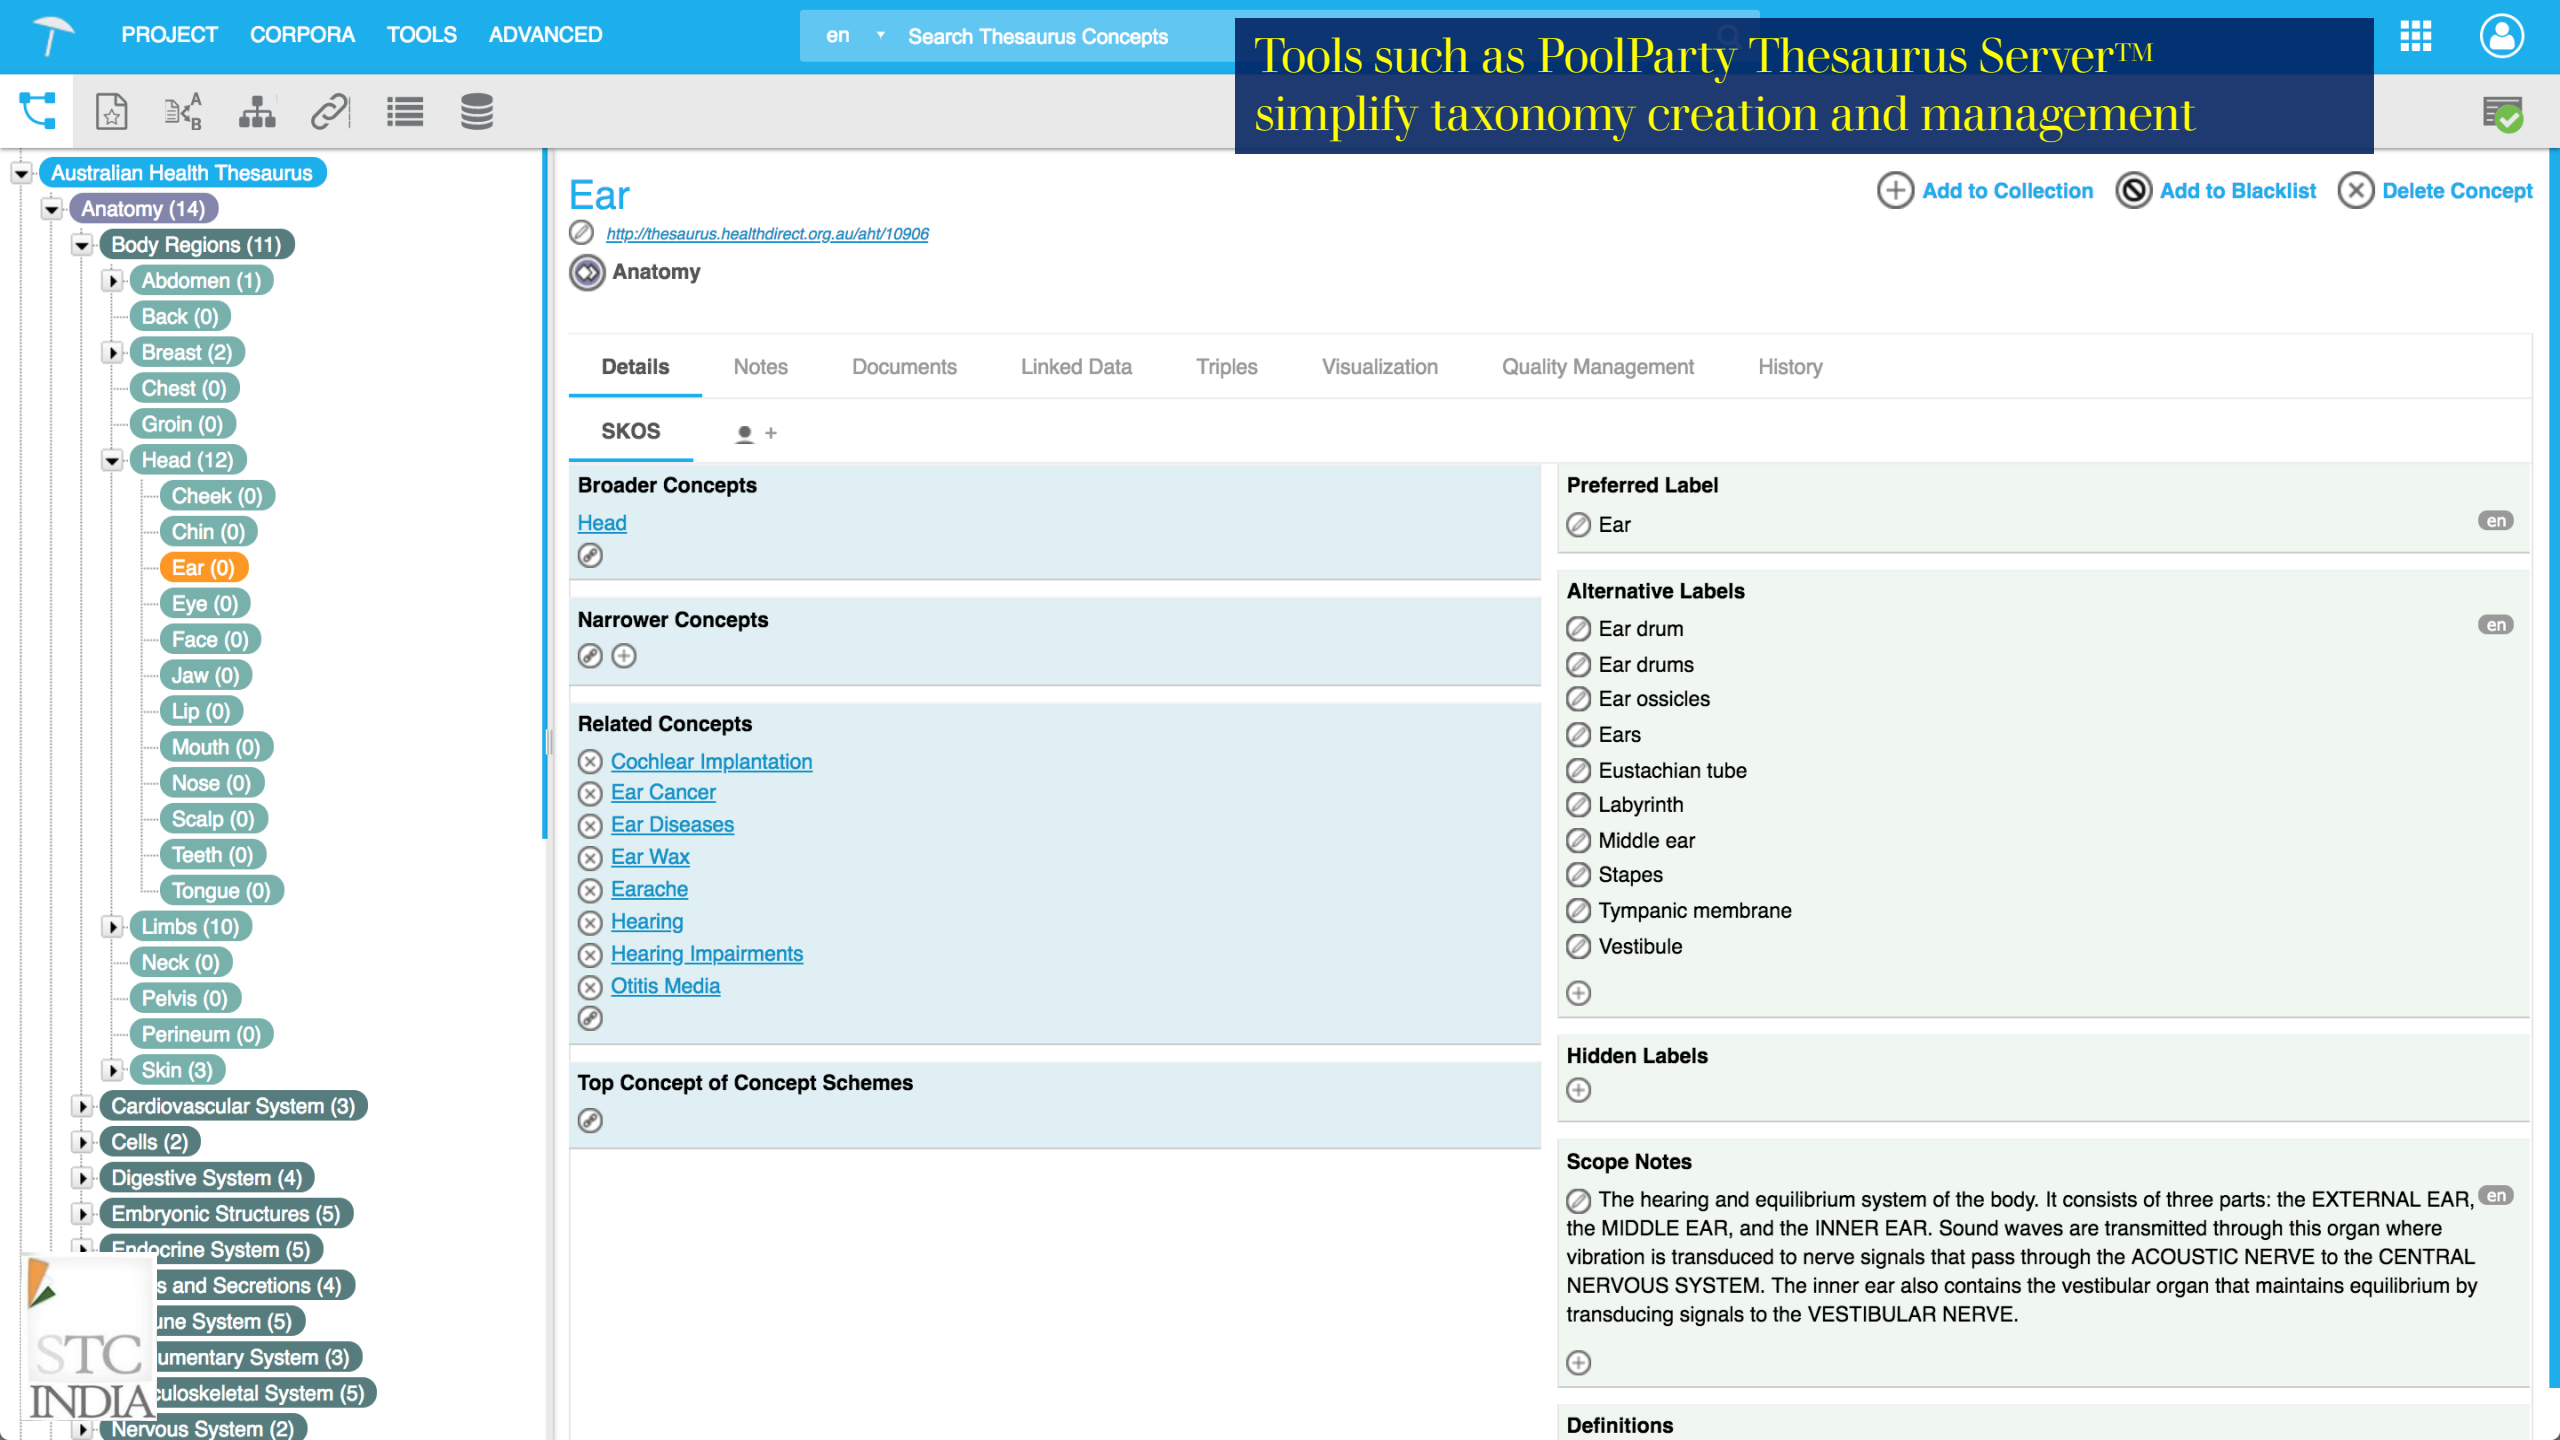
PROJECT
CORPORA
TOOLS
ADVANCED
en
Search Thesaurus Concepts
Tools such as PoolParty Thesaurus Server™
simplify taxonomy creation and management
Australian Health Thesaurus
Anatomy (14)
Body Regions (11)
Abdomen (1)
Back (0)
Breast (2)
Chest (0)
Groin (0)
Head (12)
Cheek (0)
Chin (0)
Ear (0)
Eye (0)
Face (0)
Jaw (0)
Lip (0)
Mouth (0)
Nose (0)
Scalp (0)
Teeth (0)
Tongue (0)
Limbs (10)
Neck (0)
Pelvis (0)
Perineum (0)
Skin (3)
Cardiovascular System (3)
Cells (2)
Digestive System (4)
Embryonic Structures (5)
Endocrine System (5)
Nervous System (2)
Ear
http://thesaurus.healthdirect.org.au/aht/10906
Anatomy
Add to Collection
Add to Blacklist
Delete Concept
Details
Notes
Documents
Linked Data
Triples
Visualization
Quality Management
History
SKOS
Broader Concepts
Head
Narrower Concepts
Related Concepts
Cochlear Implantation
Ear Cancer
Ear Diseases
Ear Wax
Earache
Hearing
Hearing Impairments
Otitis Media
Top Concept of Concept Schemes
Preferred Label
Ear
en
Alternative Labels
Ear drum
Ear drums
Ear ossicles
Ears
Eustachian tube
Labyrinth
Middle ear
Stapes
Tympanic membrane
Vestibule
Hidden Labels
Scope Notes
The hearing and equilibrium system of the body. It consists of three parts: the EXTERNAL EAR, the MIDDLE EAR, and the INNER EAR. Sound waves are transmitted through this organ where vibration is transduced to nerve signals that pass through the ACOUSTIC NERVE to the CENTRAL NERVOUS SYSTEM. The inner ear also contains the vestibular organ that maintains equilibrium by transducing signals to the VESTIBULAR NERVE.
Definitions

A taxonomy of viruses

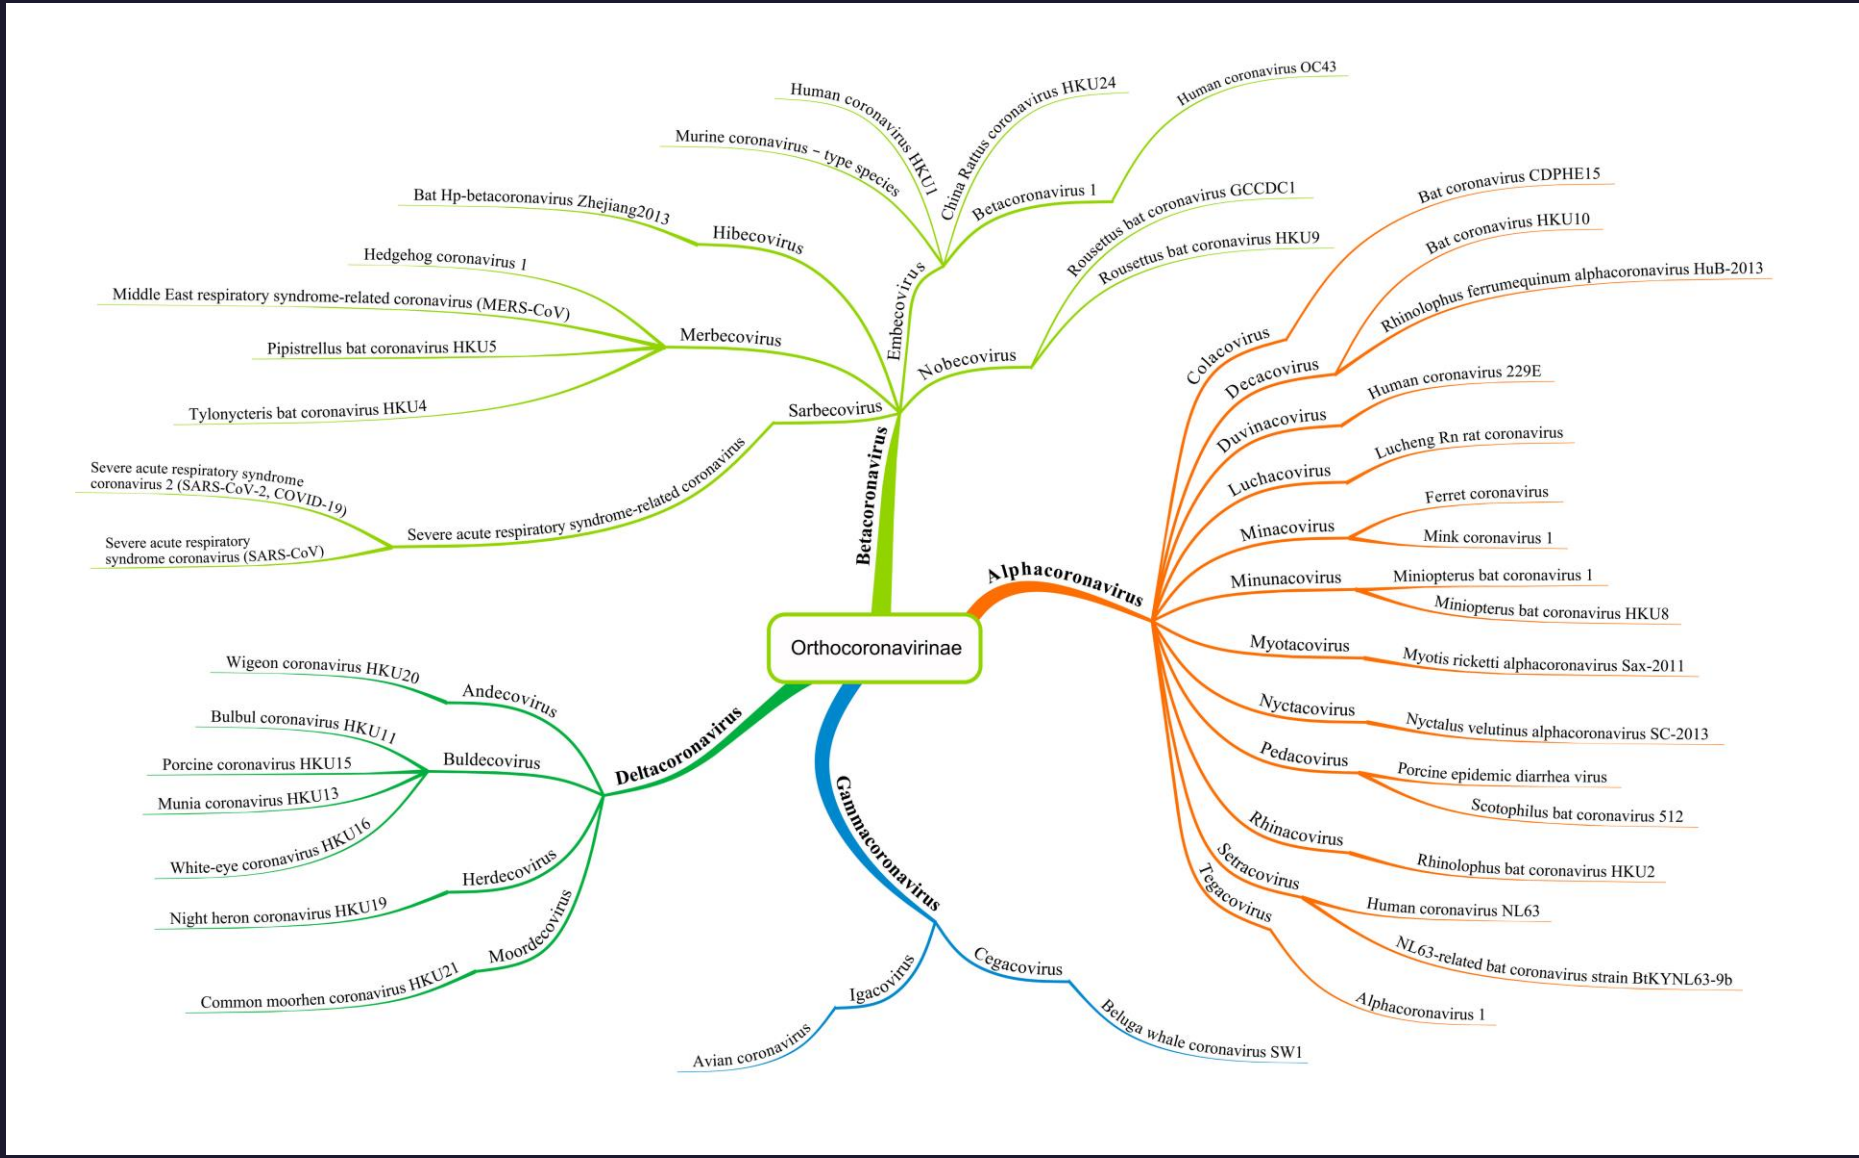

Core concept: Machine learning

- AI-driven personalization and retrieval depends upon machine learning
- It's a conversation – not limited to chatbots
- Uses natural language processing (NLP) to parse input *signals*
- Returns answers or search results based on probability
- Learns and improves with use

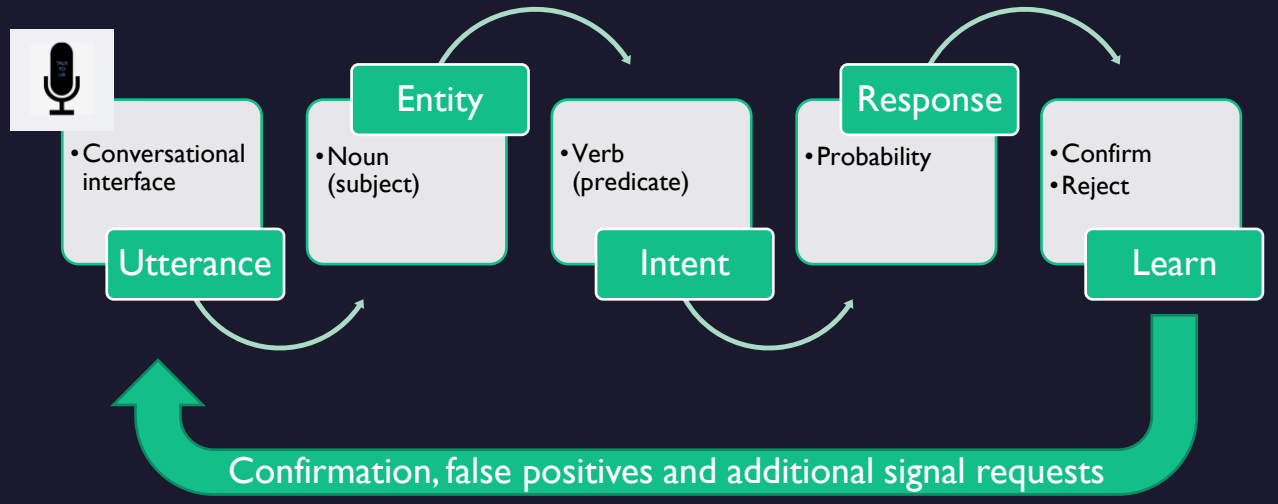

Core concept: Object container metadata

- Assign taxonomy labels to objects
  - Collection (such as a DITAMap)
  - Whole topic
  - Element (microcontent)
  - Spanned elements using <DIV> wrappers (microcontent)
- Injects explicit intent for AI discovery and retrieval
- Creates precision objects
- Harmonize these taxonomy labels with content terminology

Core concept: Object container metadata

- Assign taxonomy labels to objects
  - Collections (such as a DITAMap)
  - Whole topics
  - Elements (microcontent)
  - Spanned elements using <DIV> wrappers (microcontent)
- Injects explicit intent for AI discovery and retrieval
- Creates precision objects
- Harmonize these taxonomy labels with content terminology

```
<topic id="someunique" xml:lang="en-us">
<title>Some title</title>
<shortdesc>Some short description</shortdesc>
<prolog> <metadata>Metadata
<DATA subjecttaxo="subject label"/>
<purpose purposetaxo="purpose label"/>
<journey journeytaxo="journey phase label"/>
<experience phasetaxo="experience phase label"/>

… Plus, other taxonomy labels as appropriate (product,
feature, audience, etc.

</prolog>
<body>
<p>This is generic filler text.</p>
</body>
</topic>
```

Core concept: Autoclassification

Set-up an autoclassification service with selected taxonomy → Trainers supply representative content to the autoclassifier (ML) or a rules-based service → Trainers select the best labels (ML) → Iterate until precision exceeds 90% or higher precision → Use the autoclassifier to batch assign labels

- Automatically assign metadata to content objects by training an ML-based autoclassifier or use a rules-based service
- Ideal for classifying large volumes of existing content
- Far higher precision and lower cost than human classification

- Available autoclassification services include Watson Natural Language classifier, TopBraid EDG, and PoolParty Semantic Classifier
- Maybe we can convince the computational linguistic providers to integrate with an autoclassifier?
  It would accelerate training (crowd-sourcing).
  Acrolinx, Congree, HyperSTE – are you listening?

Core concept: The triple

- A triple is the core element of structured knowledge bases - ontologies and knowledge graphs. Both ontologies and knowledge graphs are based on "nodes"
- Consists of three basic elements: Subject + Predicate + Object
- *A triplestore* stores these models as a network of objects with materialized links between them

The triple — it is easy to build out massive networks of connected information. This structure then allows sophisticated exploration across this network and offers new insights into the organization's information. [cool video link]

Ian Piper
TELLURA INFORMATION SERVICES
From *THE KNOWLEDGE GRAPH COOKBOOK RECIPES THAT WORK*

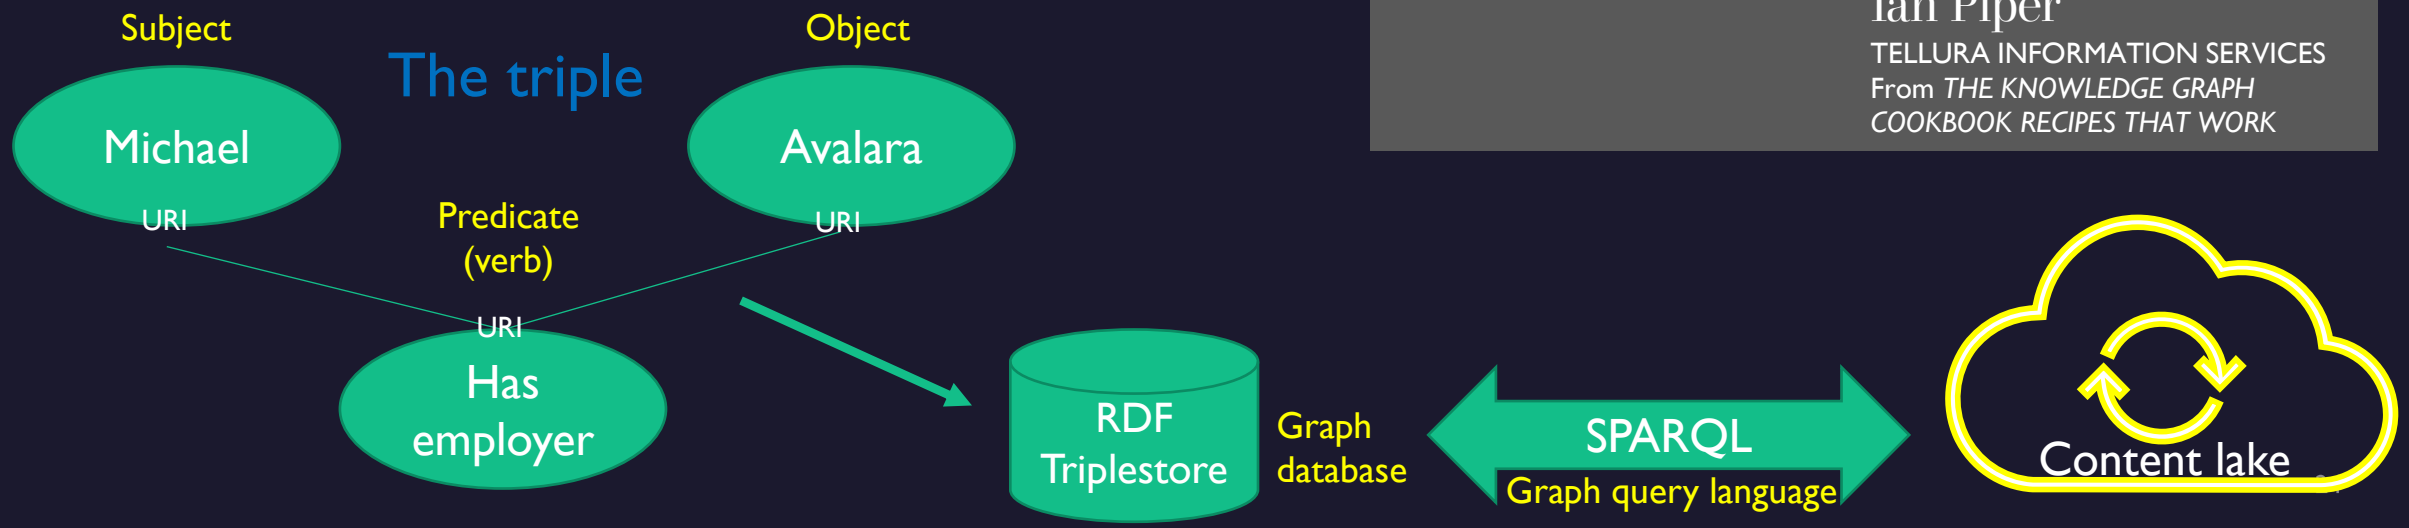

Create an ontology

- An ontology identifies *relationships* between *concepts*. It tells the machine what those relationships are for machine retrieval. It is encoded in Ontology Web Language (OWL) which is an XML schema.
- Consists of triples (Subject + Predicate + Object) and stored in a database called a triplestore (also called an RDF store) for the retrieval of triples through semantic queries.
- Acts like a schema. The DITA DTDs are a schema. It can be represented as an ontology.
- Enables terabytes of data to be reduced to only a few gigabytes of relevant data allowing more precise and effective semantic search.
- An ontology sets the foundation for a knowledge graph to capture data; it serves as the backbone for a knowledge graph (next).

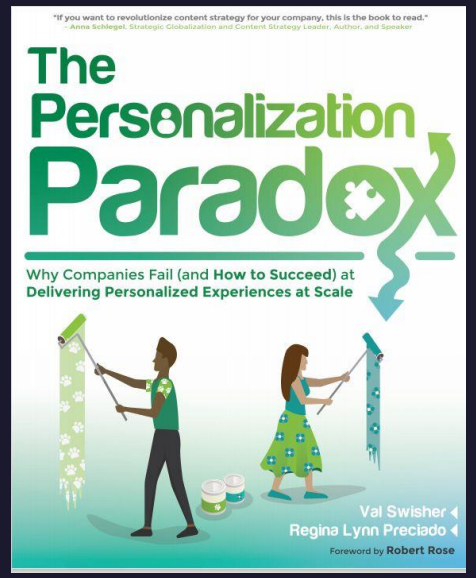

"Think of an ontology as a taxonomy of taxonomies."

Val Swisher
The Personalization Paradox

The DITA content model represented as an ontology

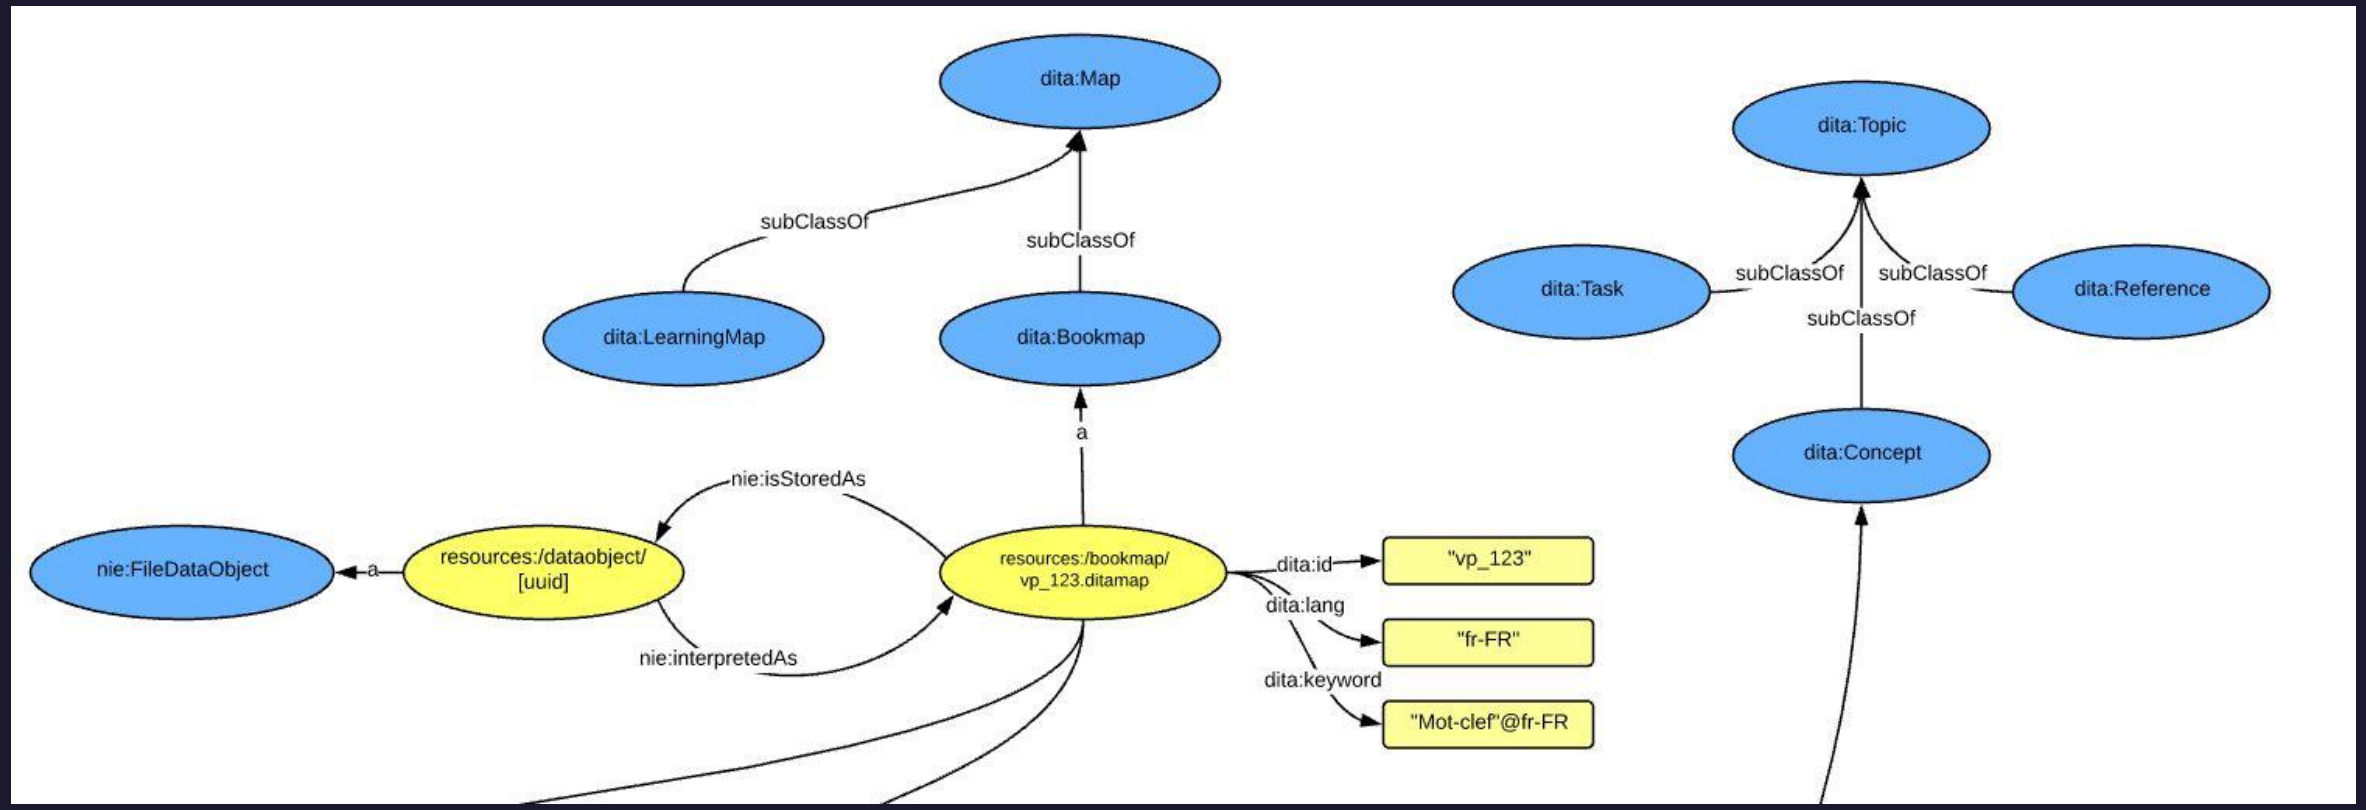

The DITA concept topic type

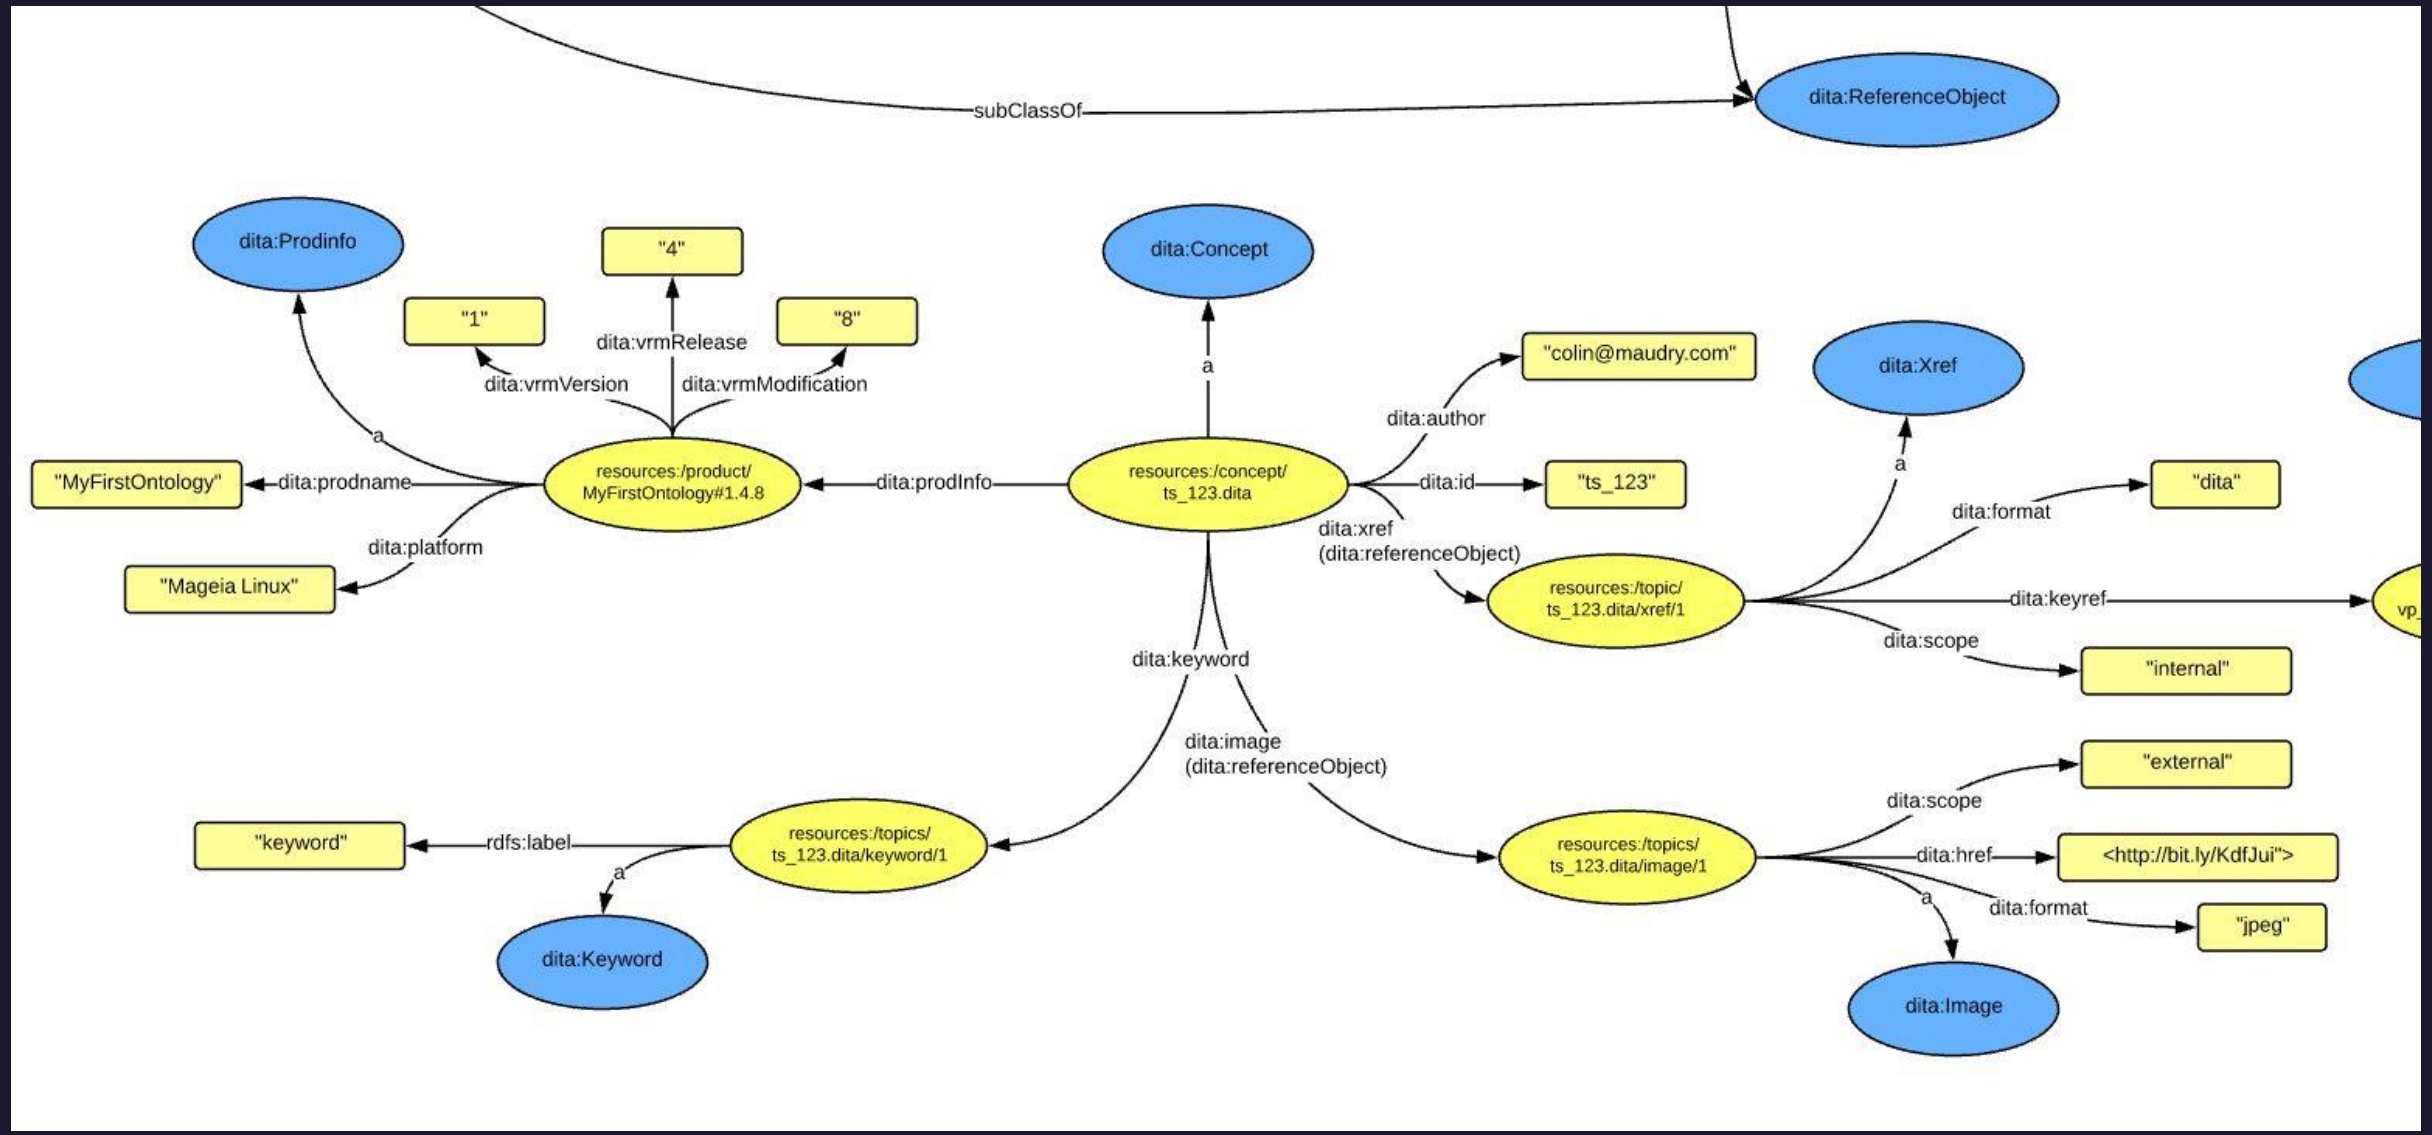

Create knowledge graphs

- A knowledge graph is a structured representation of facts, consisting of entities, relationships, and semantic descriptions captured as a semantic graph.
- Where an ontology provides a *concrete* representation of real-world objects and their relationships, a knowledge graph extends and modeled on an ontology. The graph makes extended *inferences* and *predictions;* it expands, and continually improves over time.
- The knowledge graph provides *context* and additional edge knowledge from it draws those inferences.
- Whereas an ontology specifies the formal semantics of the data, a knowledge graph captures additional intelligence over the stored data for intelligent content retrieval, organization, and delivery.

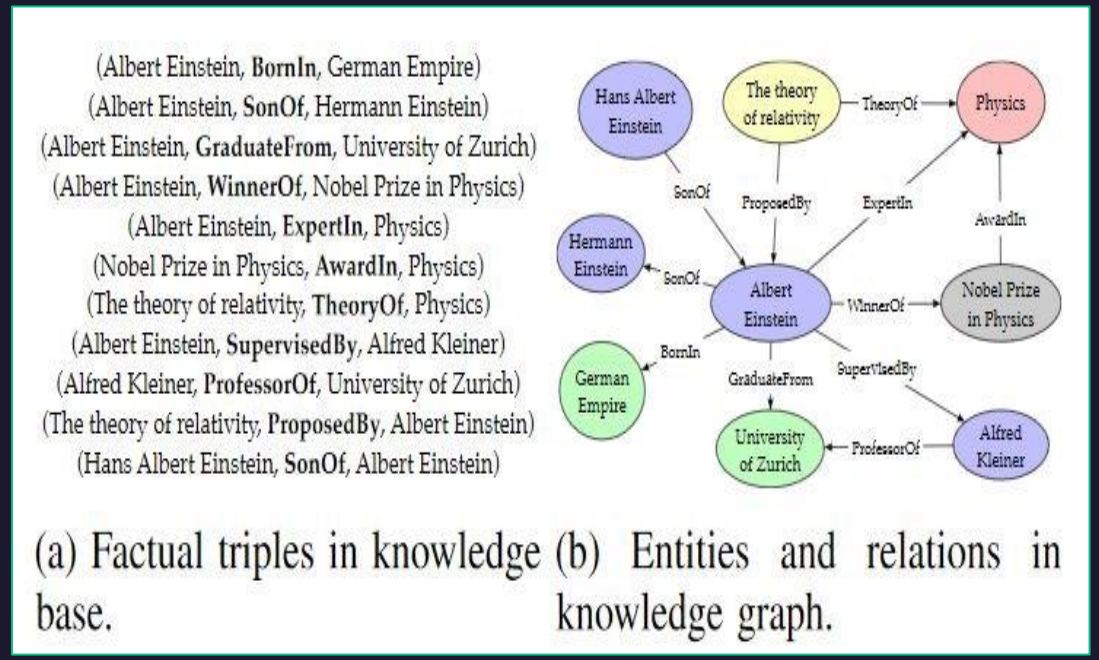

An example of knowledge base and knowledge graph

[A Survey on Knowledge Graphs: Representation, Acquisition and Applications]

Taxonomy vs. Ontology vs. Knowledge Graph

Taxonomy

- Static
- Classification of things
- Lists or hierarchical relationships
- Multi-domain
- The problem with taxonomies alone : They don't describe relationships outside of the list or across different branches of a tree

Ontology

- Describes conceptual relationships in the context of a specific domain (e.g. tax compliance)
- Ontologies only model *general* types of things that share certain properties, but doesn't include information about specific entities
- Essential to select the right combination or taxonomy and ontology for AI/ML successfully extract content for delivery
- Defines concepts and the properties that describe them
- Serves as the foundational data model for a knowledge graph. You can use an existing ontologies or develop a custom one for a specific domain.

Knowledge Graph

- An additional (virtual) data layer – for a specific instance versus a general ontology
- Extends an existing ontology; represents real-world objects and their relationships (such as content topics)
- Grows and improves over time with additional intelligence
- Links all your data together, at scale, whether structured or unstructured
- Adds extended data about individual entities (called *edge data*)
- Used by AI/ML to recognize patterns (cognitive intelligence) and predict new relationships

Building relationships for ontology and graphs

PROJECT CORPORA TOOLS ADVANCED

Search

- Ontologies
  - Cocktail_ontology
  - Families
  - FOAF
  - SKOS
  - SKOS-XL
- Custom Schemes
  - Cocktail_scheme

Create Scheme Publish

Add Relation

Definition | Description

is related — https://nextrelease.poolparty.biz/Families/is-related

Add | Person | Inverse | Person | Add

Inverse relation name — https://nextrelease.poolparty.biz/Families

Agent — No Restriction

Unique

Is a Sul

No Option

child of

Triples

Beverage (11)
Beer (0)
Brandy (6)
Fortified wine (2)
Gin (1)
Liqueur (14)
Rum (2)
Schnapps (2)
Tequila (0)
Vodka (1)
Whisky (4)

Beverages

http://vocabulary.semantic-web.at/cocktails/591cf89a-57af-49b8-9042-3fc77408c93e

Ingredients

Metadata | Cocktail scheme | Triples | Visualization | History

Subject: http://vocabulary.semantic-web.at/cocktails/591cf89a-57af-49b8-9042-3fc77408c93e

Predicate	Object	Graph
http://www.w3.org/1999/02/22-rdf-syntax-ns#type	http://www.w3.org/2004/02/skos/core#ConceptScheme	http://vocabulary.semantic-web.at/cocktails-production/thesaurus
http://purl.org/dc/terms/creator	http://localhost/user/blumauera	http://vocabulary.semantic-web.at/cocktails-production/thesaurus
http://purl.org/dc/terms/title	Beverages en	http://vocabu

Off-the-shelf solutions

- A few headless (cCMS) systems can ingest and use structured knowledge assets to create personalized experiences.
- Some can be used as a type of content middleware to *aggregate content* from multiple silos and provide an integrated enterprise content experience.
- They may use AI or business rules rather than deep learning, but they are well-positioned advanced semantic technologies.

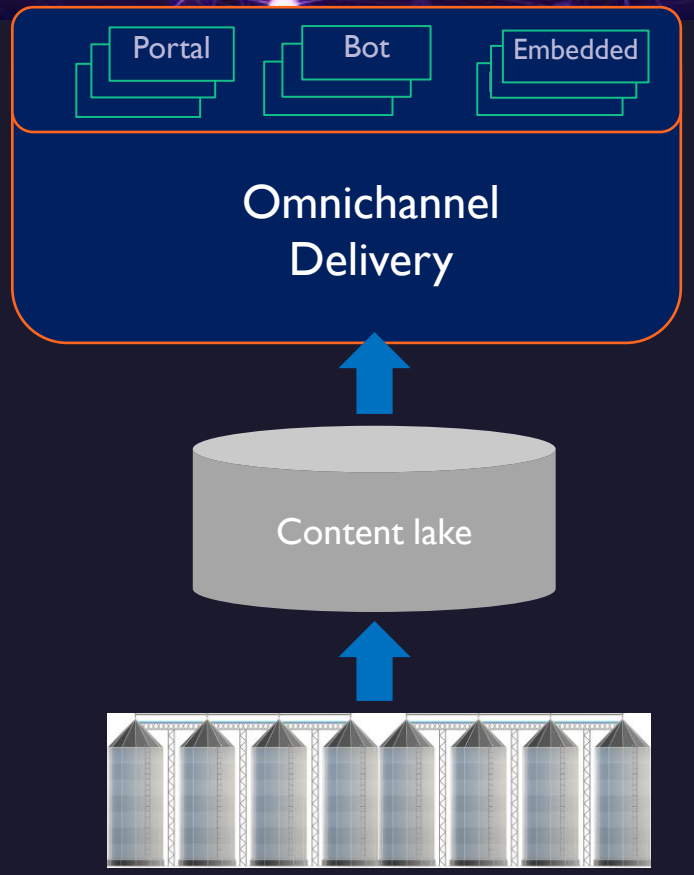

Federate content from any silo, any format

Content is functionally siloed in most organizations

- Customers forced to use multiple portals
- Disparate source formats
- Siloed production in different CMSs
- Output in multiple formats
- Siloed in isolated delivery channels
- Content cannot be integrated
  - Portal mash-ups don't solve it
  - Improved search doesn't solve it
  - Common authoring tools don't solve it

Customers want a one-stop shop for content

- They don't want to forage for content
- They don't think about content silos
- They don't want search multiple sites
- They want an *integrated* content experience

The three-tier enterprise AI /ML content supply chain architecture

content middleware (headless cCMS)

Omni-Source

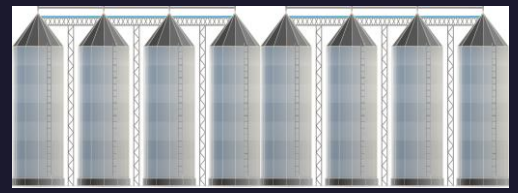

Any repository, any authoring tool, any content format or type

Content lake

AI Semantic Knowledge Assets

Omni-channel

OMNISOURCE / OMNICHANNEL CONTENT

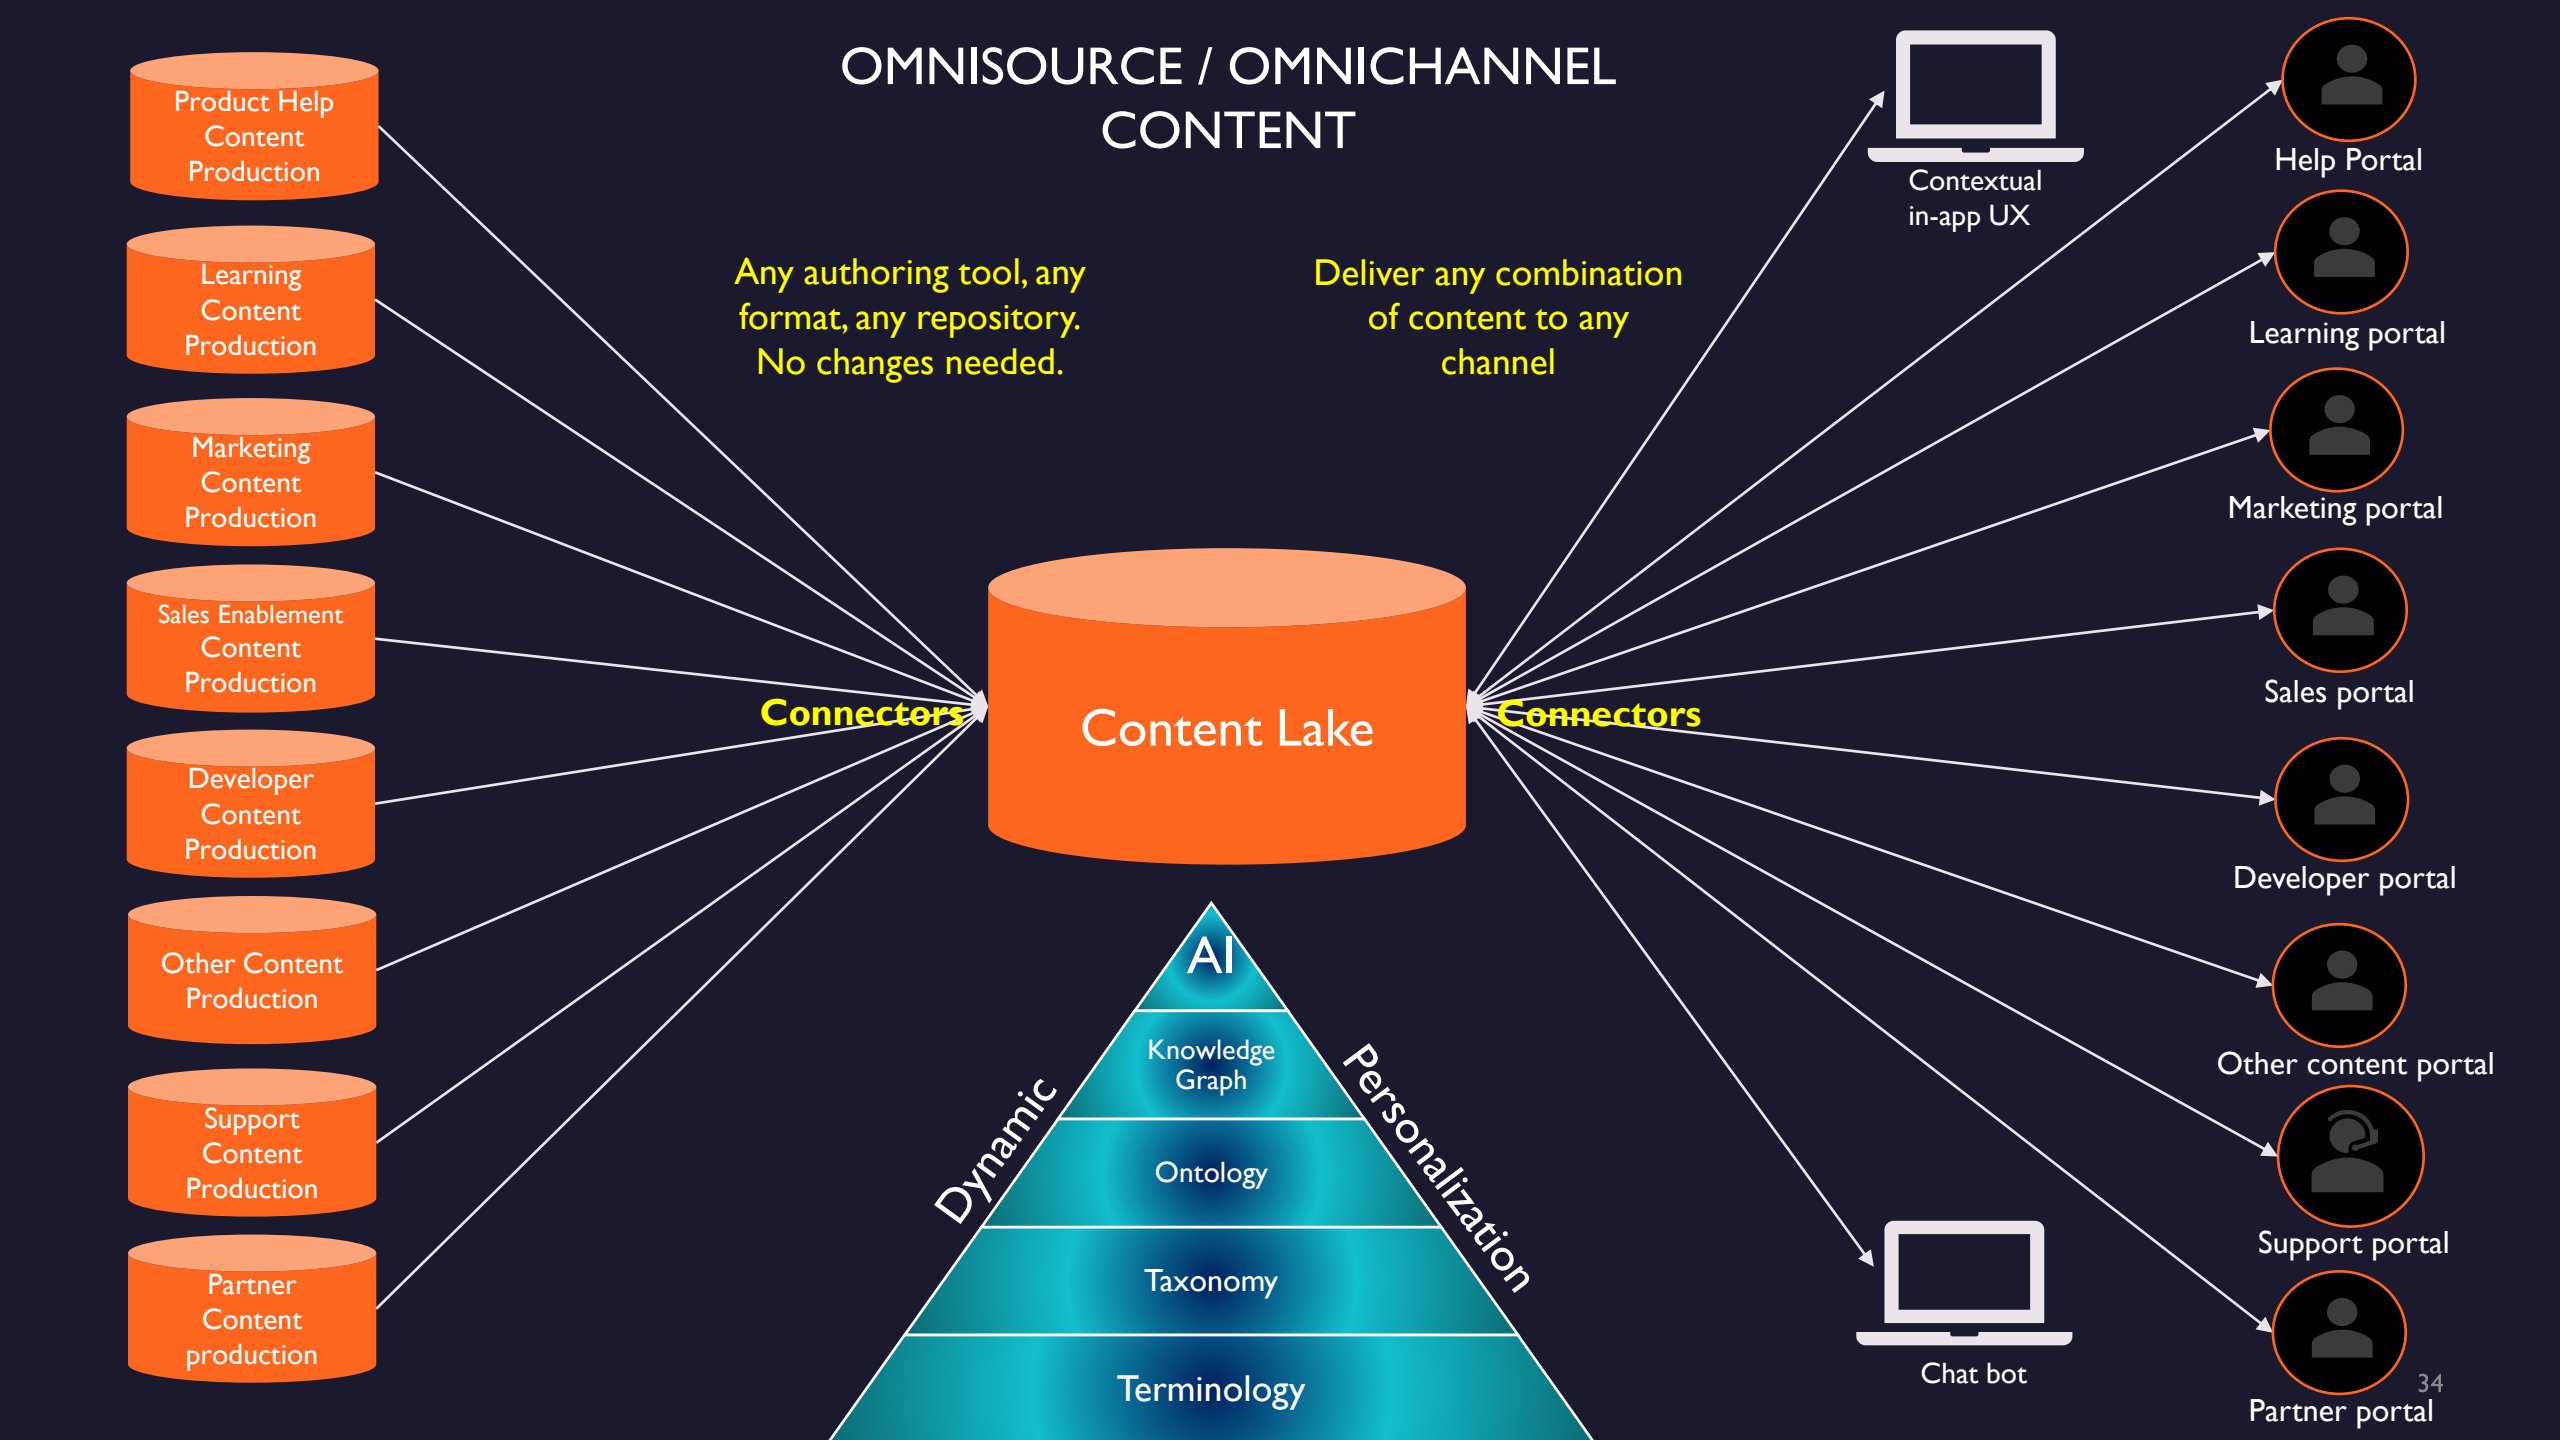

Workflow patterns and micro-workflows (scenarios)

- Businesses need to solve complex scenarios, not just one-off tasks. Most are sequenced, multi-task workflows.
- We're used to constructing prescriptive collections and navigation that support everyone, but no one in particular.
- We can model and encode prescribed product use patterns to do dynamic construction of navigation for machine retrieval of multiple objects to support multi-task scenarios?
- *Workflow patterns* and *micro-workflows* can be captured in machine-readable format and used as initial patterns for dynamic content organization with a knowledge graph.
- He also become the basis for machine-learning based on what the user does with the resulting dynamic collections.

Moonshot
Weaving it all together

- Generate multiple graphs from multiple sources, including
  - Domain graph from content corpus or corpora from semantically self-describing and enriched (taxonomy) objects.
  - Graphs from user personalization metadata
  - Graphs from patterns for automatic content organization
    - Workflow patterns and micro-workflows
    - TOCs (can be derived from DITAMaps)
    - …even relationship tables!
- Use linked or stacked graphs to mine one another
- Mine the graphs for insights not previously feasible or possible
- Generate a result graph (or map) – then feed it to a cCMS for dynamic search, retrieval, organization, and delivery

Achieving proactive, dynamic, non-prescriptive search, retrieval, organization, and delivery

A whole new genre of graph-powered content applications

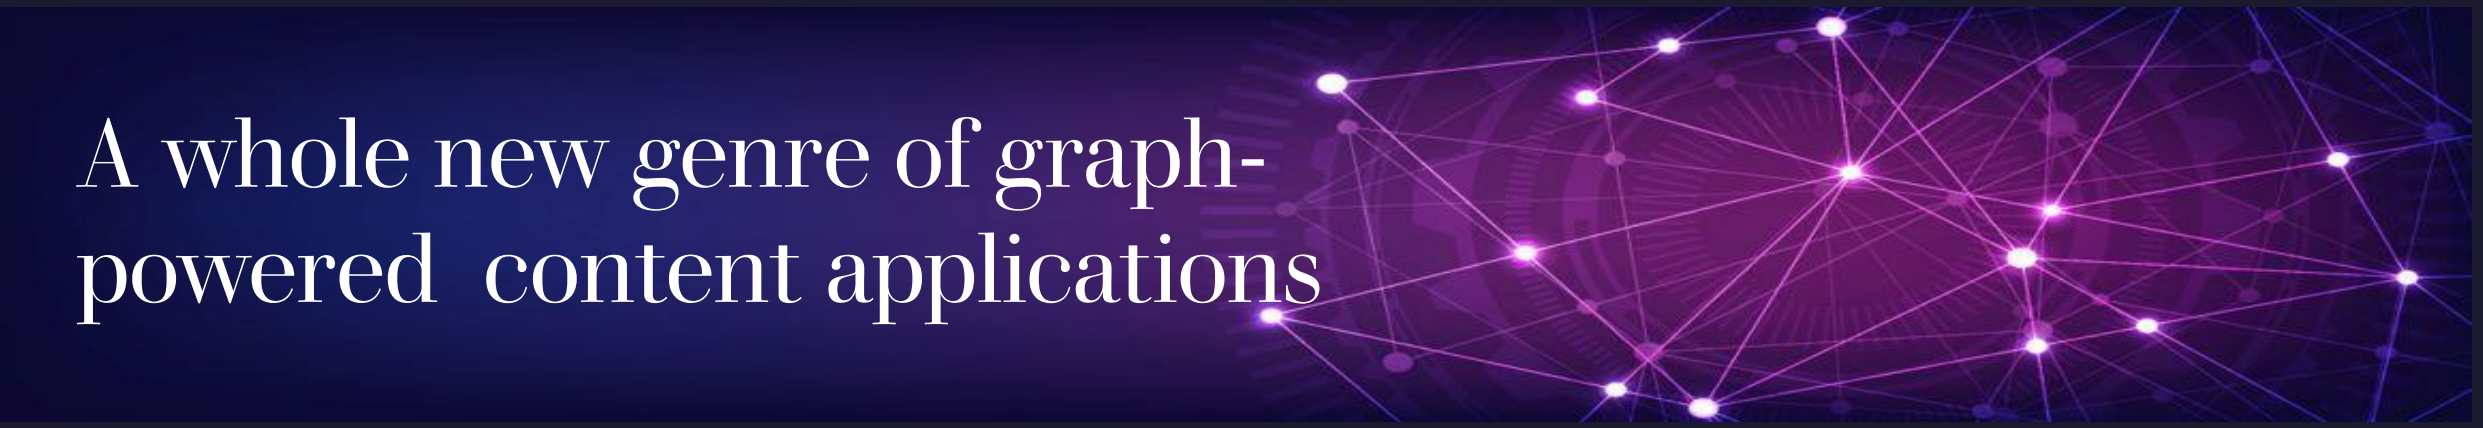

Some graph-driven application classes:

- Discovery and insights
- Precision answers (chatbots)
- Smart recommenders
- Dynamic assistance
- Autonomic content

- Mine a graph of your content corpus to discover duplicate, overlapping, or conflicting information
- Discover the effectiveness of content, not only what content is used, and which content is unused or underused, but why and what content do they use instead, or what alternate actions do the users take
- Answers to previously elusive questions such as what role post-sales content plays into revenue generation, conversions, upsell and cross-sell.

- A dynamic collection and its sequence can be altered based on the path a user takes in real-time.
- Models can self-adjust based on what the user does with the dynamic content, such as save and reconfigure personalized content collections.
- Progressive content can be truly dynamic and interactive. As users request more information the content becomes *elastic*; it can expand or contract on the fly. Moreover, the system can learn from the inbound user signals, reconfigure, and heal itself.

Dynamic content and dynamic content for scenarios

Dynamic content using only taxonomy

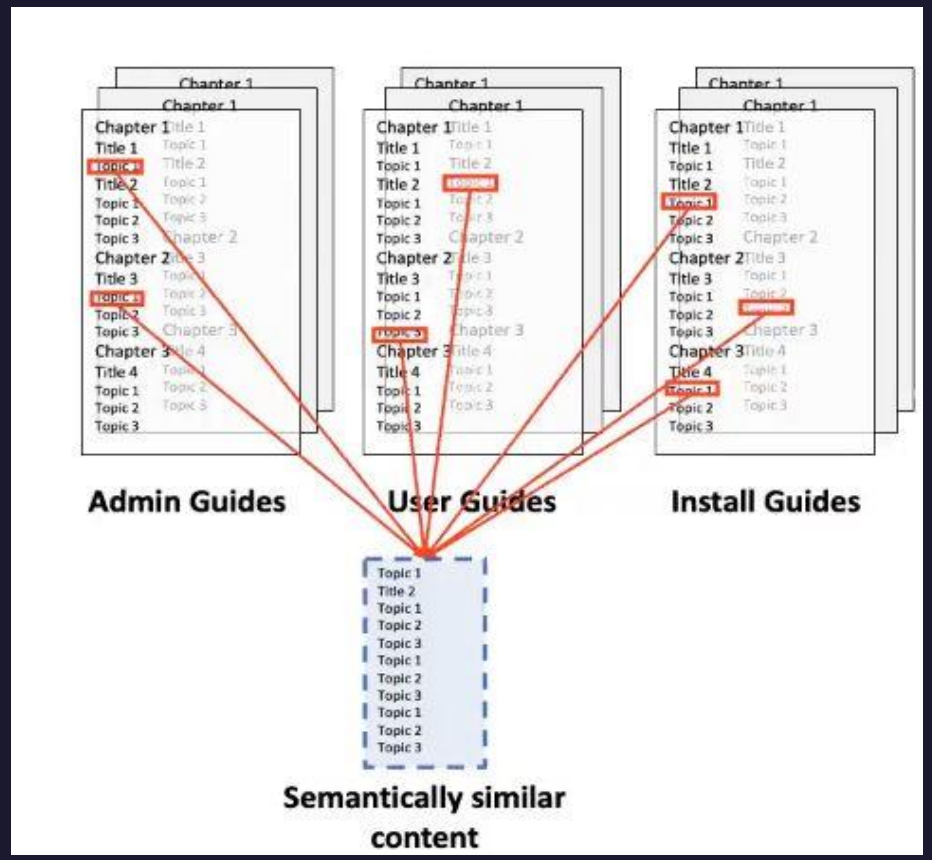

Graph query + workflow pattern

Automatic machine-sequenced table of contents to assist with complex scenarios

Autonomic content

- Responsive content – automatically change, expand, contract, or repair content.
- Use the constantly changing signals to adjust the content scope and configuration as the cognitive system learns – even make automatic corrections.
- Natural language generation (NLG)[1]
- Robotic content generation, such as Videate™

- Intelligent (interactive) highlighting
  - During content creation: AI training for defining micro documents (elements and spanned <DIV> content elements)
  - End-user highlighting: Crowdsource adjustment of content object scope
  - AI/ML learning based on custom collection assemblies

[1] Wordsmith, Automated Insights

Our roadmap

Create intelligent objects

Align content models

Active terminology governance

Objects (DOM) | Content types | Content models | Content journeys | Terminology

Define and harmonize content types

Define and map macro *and* micro content Journeys

Create key taxonomies

Create relational schema

Ingest and use with semantic application(s)

Taxonomy | Autoclassify | Ontology | Knowledge Graphs | Deploy

Assign labels to content objects

Enrich with intelligence, then mine them

A glimpse of the future

What else might we imagine...

(BHAGs – Big, Hairy, Audacious Goals)

Hyper-personalization

- User signals (utterances) are not static! Yet that is often the default assumption (that is, wait for the next query). We cannot not communicate! We can enable our systems for *continuous signal acquisition*.
- Utterances can be captured as structured content rather than queries and mapped against KGs?
- What if we could virtually watch over the user's shoulder (with permission of course)? "Help needed on aisle 6!"
- Including emotive and other signals.
- Blockchain just might provide that ability…

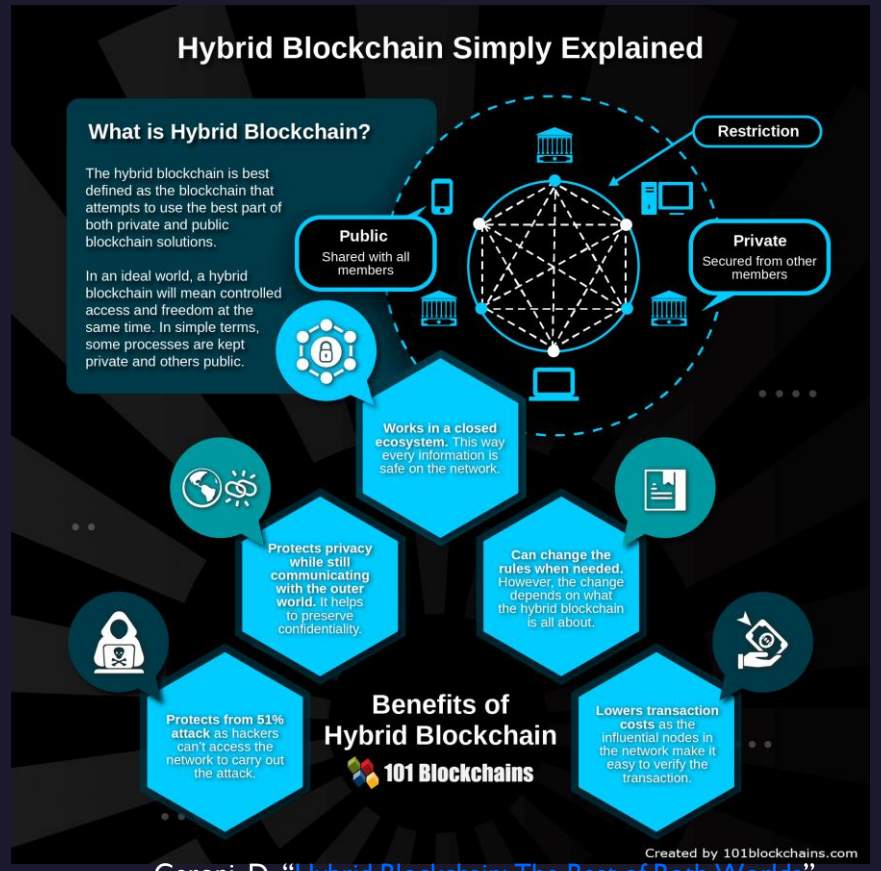

Geroni, D. "Hybrid Blockchain: The Best of Both Worlds", 101 Blockchains, January 28, 2021

Blockchain for responsive and adaptive user assistance

- Blockchain isn't just for Bitcoin or financial transactions, It's an immutable ledger of transactions, which to us equate to utterances
- Each user assigned their own wallet – continuous record of transactions over time
- Public-private model – anyone can participate, but impossible to hack
- Improves signal-to-noise ratio
- Microcredits: Content objects as NFT's to integrate and deliver 3rd-party content by merging content knowledge graphs!

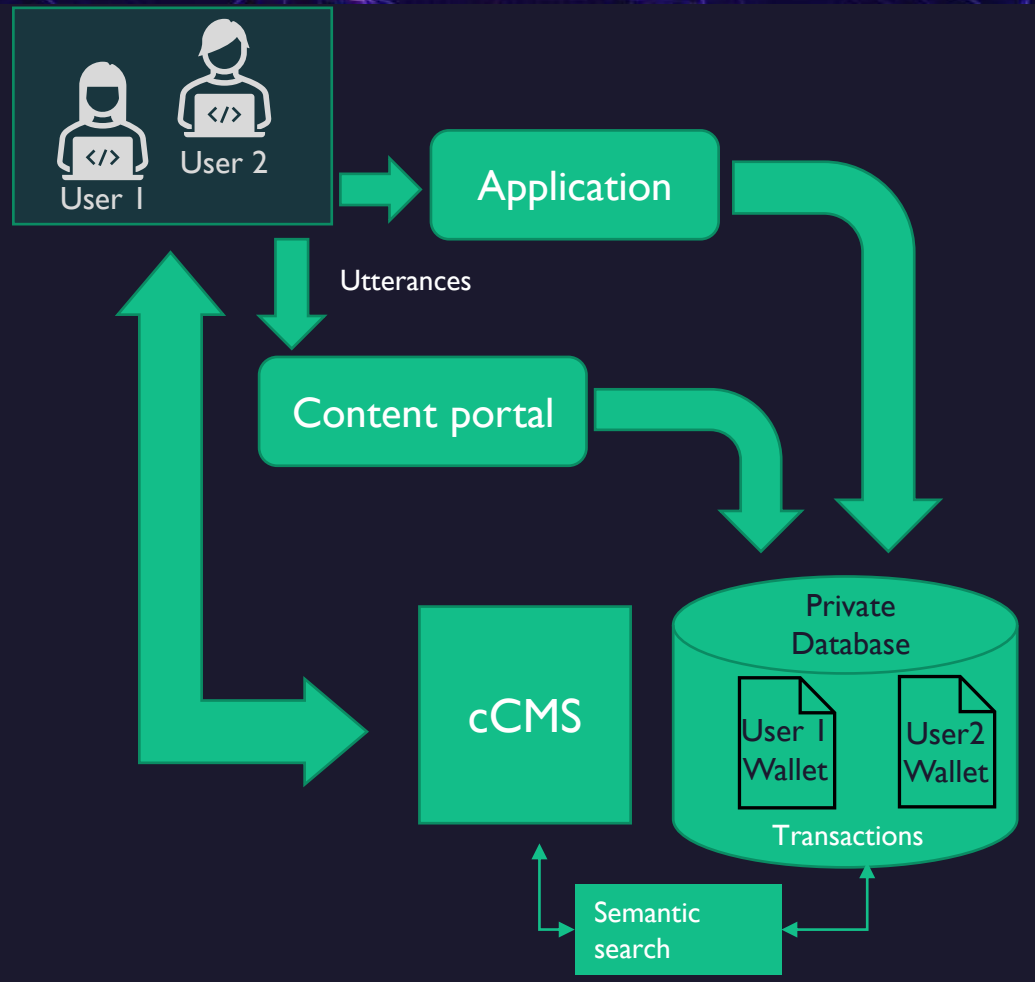

Summary

Summary

How intelligent is your content? Or more precisely, how intelligent is your content supply chain? Can your supply chain learn?

A cognitive content supply chain can provide dynamic one-on-one personalization that is built upon an intelligent content supply chain that includes content *delivery* intelligence – one that combines taxonomy, ontology, knowledge graphs, and cognitive AI services to achieve adaptive, pro-active, and autonomic content with a single intelligent content architecture.

Download full paper: *ThinkingDocumentation.com*

Semantic content graph guild (public discussion forum): https://thinkingdocs.com

References

[1] Iantosca, M. From Microcontent to Neurons: A practical guide for building a cognitive AI content supply chain for highly personalized user assistance, 2021

[2] Swisher, V, Preciado,R.L. The Personalization Paradox, XML Press, 2021

[3] Ji, S., Pan, S., Cambria, E., Marttinen, P., and Yu, P. S., A Survey on Knowledge Graphs: Representation, Acquisition and Applications, *arXiv e-prints*, 2020.

[4] Blumauer, A., Nagy, H. The Knowledge Graph Cookbook: Recipes That Work, *monochrom*, 2020, *Semantic Web Company*, 2020.

[5] Geroni, D. Hybrid Blockchain: The Best of Both Worlds, 101 Blockchains, January 28, 2021

[6] Manhaes, M., Ko, T., Selim, A., Amer, 0., Sri, L. Building Cognitive Applications with IBM Watson Services: Volume 4 Natural Language Classifier, 1985 IBM Redbooks

[7] Kotov, Alexander. (2017). Knowledge Graph Entity Representation and Retrieval.

[8] Park, Y., JungHyen, A. Methodology for Automatic Ontology Generation Using Database Schema Information, 2018 *Advances in Mobile Networking for IoT Leading the 4th Industrial Revolution*

[9] *George, Palliyathu Vishal (Babusapalya, IN), Iantosca, Michael J. (Wake Forest, NC, US), Kurian, John (Bangalore, IN), Sankar, Balaji (Bangalore, IN), 2019 UNSTRUCTURED DOCUMENT MIGRATION, United States International Business Machines Patent 20190205460*

[10] Palliyathu, Vishal; Tiwari, Deepti; Kothari, Anu; Upendra, Vineet; and Gupta, Ritu, DYNAMIC GENERATION OF AN ONTOLOGY-BASED AI SCHEMA FOR CHATBOTS, Technical Disclosure Commons, (June 29, 2020)

[11] Dysart, j., The Robots Cometh: How Artificial Intelligence is Automating Writing Jobs, 2021 *Robot Writers AI*

[12] Hedden, H., The Accidental Taxonomist, 2010 *Information Today, Inc.*

[13] Jha, Ajit K Artificial Intelligence for Business Leaders: ARTIFICIAL INTELLIGENCE and MACHINE LEARNING BOOK FOR MANAGERS, LEADERS, 2020 Independent

[14] Bahareh Sarrafzadeh and Olga Vechtomova. 2014. Combining document retrieval with knowledge graphs for exploratory search. In *Proceedings of the 5th Information Interaction in Context Symposium* (*IIiX '14*). Association for Computing Machinery, New York, NY, USA, 345–347. DOI:https://doi.org/10.1145/2637002.2637060

[15] Bahareh Sarrafzadeh, Adam Roegiest, and Edward Lank. 2020. Hierarchical Knowledge Graphs: A Novel Information Representation for Exploratory Search Tasks. ACM Transactions on Information Systems 4, TOIS, Article 1 (April 2020), 41 pages.