Fiesole Collection Development

The Retreat

2 April 2019

Information, Intelligent Machines, and New Knowledge

Ruggero Gramatica, PhD

Yewno

Epistemology.....

Is the theory of **Knowledge**.....

its methods....

its scope....

..... and the rationality of beliefs.

[source: Oxford Dictionary]

Knowledge....

… is the understanding of facts, information, descriptions, skills

… acquired through experience, education

…or by perceiving, DISCOVERING

… or LEARNING

[source: Wikipedia]

SEARCH

- When you know what you're looking for…

Ask "What"

Simple questions → Simple list of results

DISCOVER

- When you don't know what you don't know

Ask "Why"

Thoughtful questions → New discoveries

The World Economic Forum Artificial Intelligence Panel, Davos 2016

"A [traditional] search engine is incredibly good at processing and indexing documents and sometimes returning useful ones when we put in a good query, but they understand little or nothing of the content of the document so they can't really answer your question, they may be giving you back a document that contains the wrong answer…"

…Whereas if these systems can really understand everything that the documents contain, at least in a factual sense, then they can be far more useful. If the search engine industry is worth a trillion dollars right now then this new technology could be worth 10 trillion because it will have so many more applications and be so much more useful to so many people…"

OVERVIEW

- To present a cognitive framework that unifies different areas:

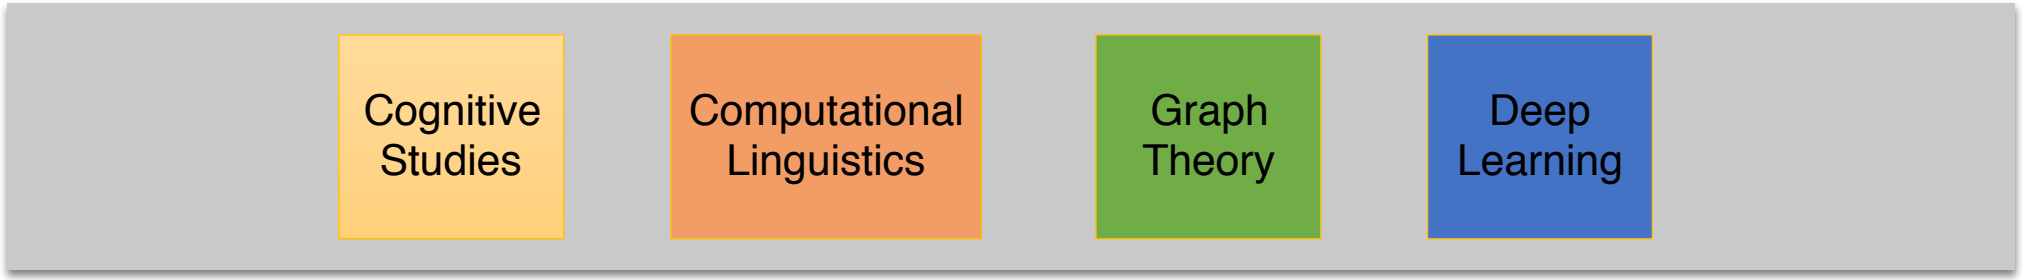

- We want to be an Anthawk, an imaginary species half Hawk and half Ant.

Trying to get a view of the entire landscape, but also being able to reach down to the empirical ground and grasp all the interesting details.

THE PROBLEM…

IS IT ONLY ABOUT TOO MUCH INFORMATION?

- The volume of information accessible today is constantly growing and appears fragmented and dispersed through a multitude of heterogeneous sources.

- It is virtually impossible for an individual to access all the information components made available by various sources and process the content

NOT ONLY: INFORMATION IS NOT THE SAME AS KNOWLEDGE!

- Having access to information is NOT the same as knowing

- Every discipline or subject requires information access and processing so as to synthesize its contents and form Knowledge.

?

But how can we access ALL the information spectrum available and derive knowledge of a topic?

“

- What distinguishes Information from Knowledge is the way that knowledge empowers the intellectual and physical capabilities of individuals.
- Knowledge is a matter of cognitive capacity that creates an active ability to think and reflect.
- Information alone is a passive cognitive process, without any special utility if not transformed into knowledge.
- Finally, Knowledge provides the means by which information is interpreted and brought to life.

”

– The Work Foundation’s Knowledge Economy Programme interim report (Brinkley 2008)

From Information

to

Extraction of Knowledge

Yewno

THE RENAISSANCE OF ARTIFICIAL INTELLIGENCE

- Artificial Intelligence (AI) was born in the 1950s as an attempt to build models capable of autonomously analyzing complex systems and was inspired by the progress made in studying the structure of the human brain

- For decades, up to the beginning of the new millennium, models based on simplified neural networks required enormous low-scalable IT infrastructures with large amounts of data to be analyzed and were inefficient in simulating solving real problems

- This period is known as "the A.I. - Winter ", is now replaced by the Spring of Artificial Intelligence thanks to progress made in the study of cognitive processes, Machine Learning-based models, Neural Networks, and - more recently - Deep Learning

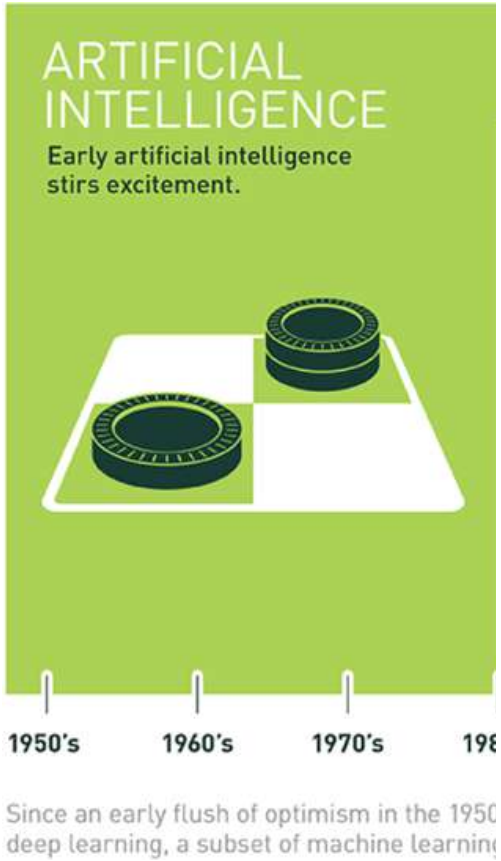

A.I. AND NEURAL NETWORKS - WHY NOW?

Artificial intelligence (AI) has recently undergone a renaissance, making major progress in key domains such as vision, language, control, and decision-making.

This has been due, in part, to today's cheaper computational resources, which aligns with the natural strengths of deep learning.

Simply put, a Neural Network is a:

Mathematical Model

structured like a human brain and able to simulate some of its functions.

A.I. AND NEURAL NETWORKS

The structure of a Neural Network

It is composed of several computational nodes called Neurons, interconnected by a layer structure

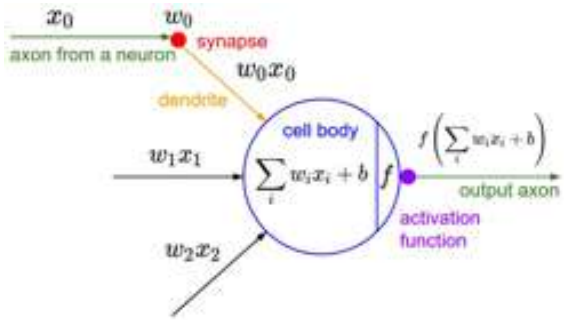

A Neural Network with enough number of nodes and layers can understand the characteristics of a complex (nonlinear) system and provide guidance on the dynamics of its constituent elements

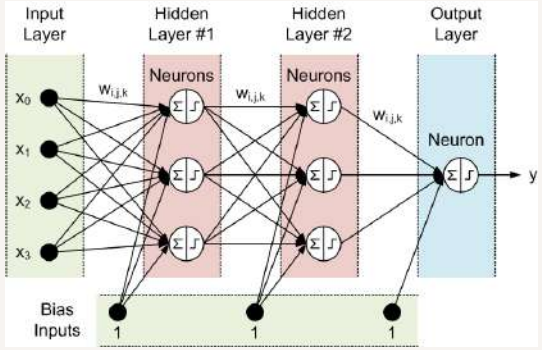

WHAT IS KNOWLEDGE?

- It is to understand the information spectrum of a topic.
- It is the condition at the outset of learning the truth of a fact through reasoning.

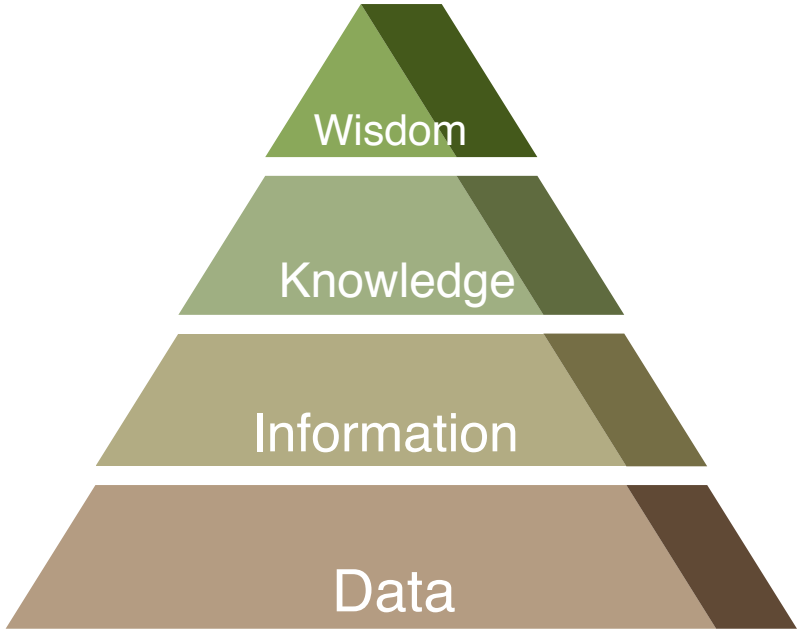

Knowledge is the interpretation of the Information

Knowledge is formed in our mind

Many of the human cognitive abilities are articulated from Knowledge Acquired in Unity, during our cognitive formation:

- How do we orient ourselves through the surrounding world?
- How do we solve problems from the simplest ones to the more complex ones?
- How do we understand meaning, make deductions and decisions?

The problem of Knowledge representation

HOW CAN WE REPRESENT KNOWLEDGE?

A new discipline: Computational Cognitive Linguistics

How language and symbology forms our thoughts and Knowledge is gained.

Thanks to the seminal work of:

Peter Gärdenfors and Ray Jakendoff

Peter Gärdenfors

Professor of cognitive science
University of Lund, Sweden

Ray Jakendoff

Professor of philosophy
Seth Merrin Chair in the Humanities

Director of the Center for Cognitive Studies, Tufts University

- According to the cognitive tradition, meanings are mental entities. The core idea of cognitive linguistics is that meanings of linguistic expressions and other communicative acts are mental entities.

- We pursued the construction of a framework based on the assumption that our minds organize the information through cognitive processes and in a format that can be modeled in geometric or topological terms. The conceptual spaces.

CONCEPTUAL ABSTRACTIONS

- Concepts are mental abstractions

  based on human cognitive functions, including learning and reasoning

- Recent studies at the Computational Neuroscience Gallant Lab at the University of California, Berkeley - show how our brains can dynamically map concepts and categories

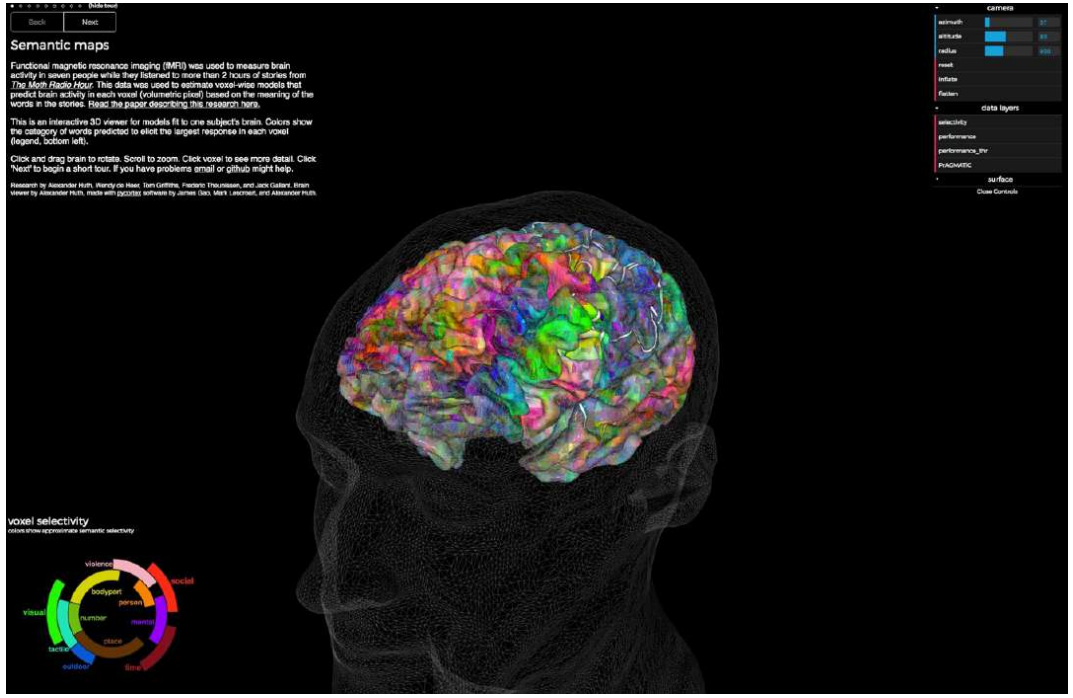

Decoding the Semantic Content of Natural Movies from Human Brain Activity
A.G. Hush, et all - Front. Syst. Neurosci., 07 October 2016

The Gallant Lab at UC Berkeley - Brain Viewer - http://gallantlab.org/index.php/brain-viewer/

CONCEPTUAL ABSTRACTIONS

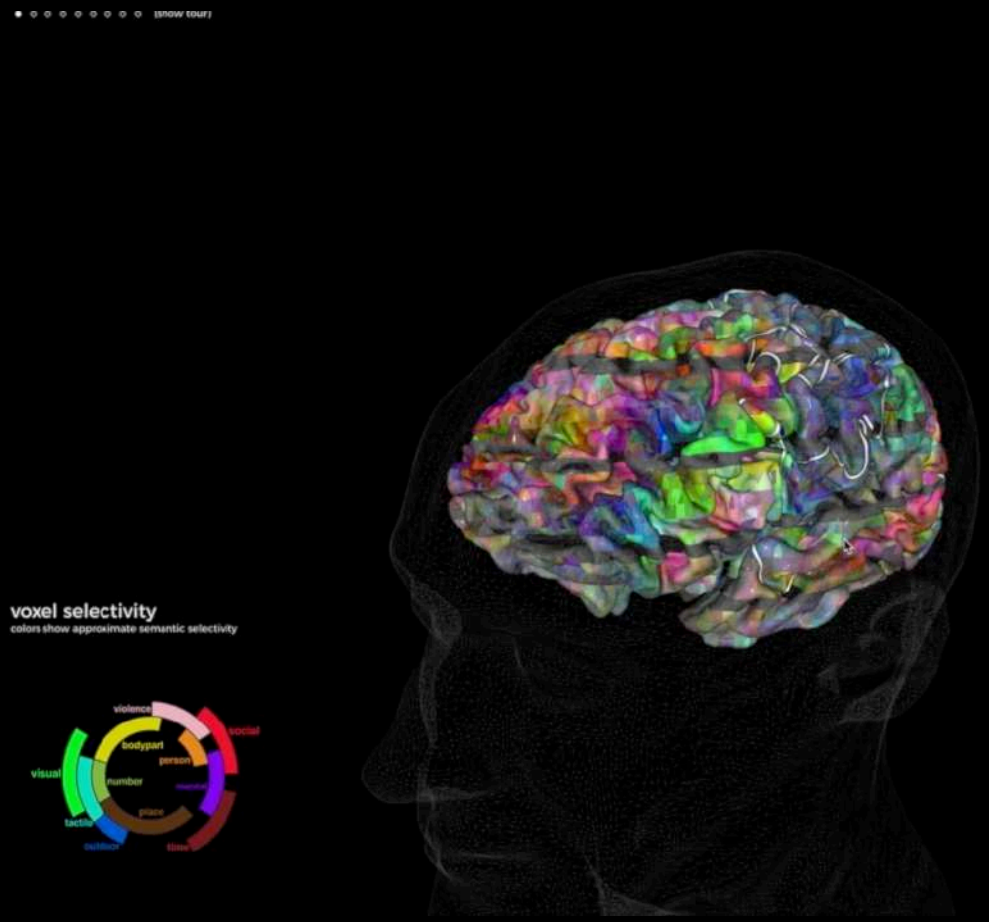

The Gallant Lab at UC Berkeley - Brain Viewer - http://gallantlab.org/index.php/brain-viewer/

REAL WORLD VS. PROJECTED WORLD [JACKENDOFF (1983)]

We have a conscious access of the real world only through the projected world…

The world is unconsciously organized by our mind

And in fact we can talk about things only if they have achieved a mental representation through these processes of organization.

A projection of the world is in our mind

CONCEPTUAL SPACES

- And here is our intuition:

What if we can geometrically represent a concept?

How can we project it into a multidimensional space?

- The idea is that the meanings that we assimilate can be described as organized in abstract spatial structures that are expressed in terms of dimensions, distances, regions, and other geometric or inferential notions.

(Gärdenfors: Conceptual Spaces, 2000 - The Geometry of Meaning, 2014)

CONCEPTUAL SPACES

What is a concept?

A Concept is an abstraction of

- an idea,
- a thought, or
- an expression

portrayed in various forms

We can "design" a Concept as an atomic information unit consisting of:

- one or more definitions
- belonging to one or more categories/topic

SINGLE WORD

Democracy (from ancient Greek: δῆμος, démos, "people" and κράτος, krátos, "power") etymologically means "government of the people", a system of government in which sovereignty is exercised, directly or indirectly, by all Citizens who resort to a vote.

SHORT PHRASE

Quantitative Easing is a monetary policy tool, and it designates one of the ways in which a central bank creates coinage and its injection, with open market operations, in the financial and economic system.

LONG PHRASE

The Decline of Roman Empire was the process of decline in the Western Roman Empire in which the Empire failed to enforce its rule, and its vast territory was divided into several successor polities.

INTELLIGENT FRAMEWORK [R. GRAMATICA et alii - 2014, 2017, 2018]

Let's look at an intelligent framework capable of:

- Reading the structure of data sources and understanding meaning
- Identifying concepts, definitions and distinguishing the differences
- Identifying concepts, definitions and distinguishing the differences

World War II

World War II, also called Second World War , conflict that involved virtually every part of the world during the years 1939–45. The principal belligerents were the Axis powers …

Illuminism

The Enlightenment (also known as the Age of Enlightenment or the Age of Reason) was an intellectual and philosophical movement that dominated the world of ideas in …

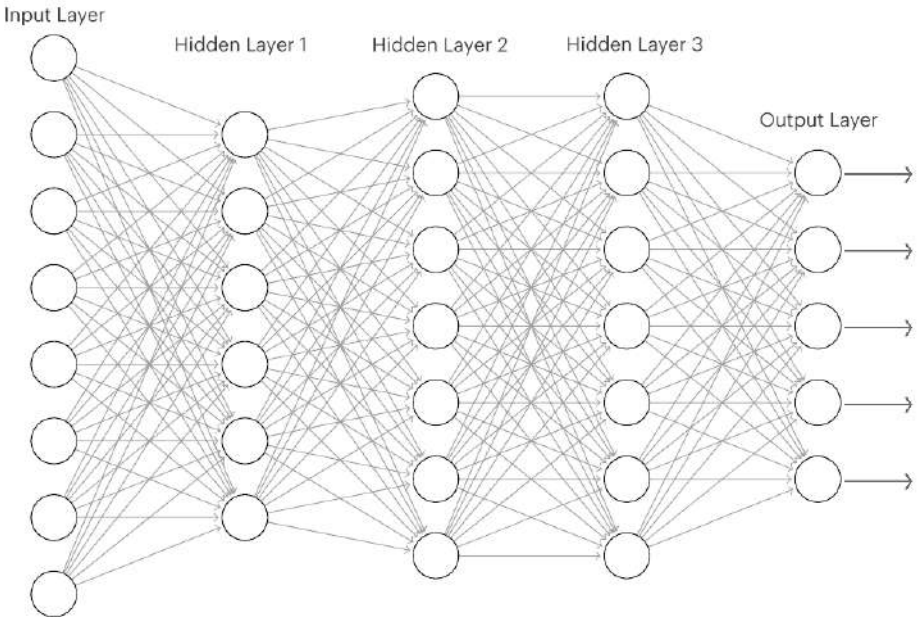

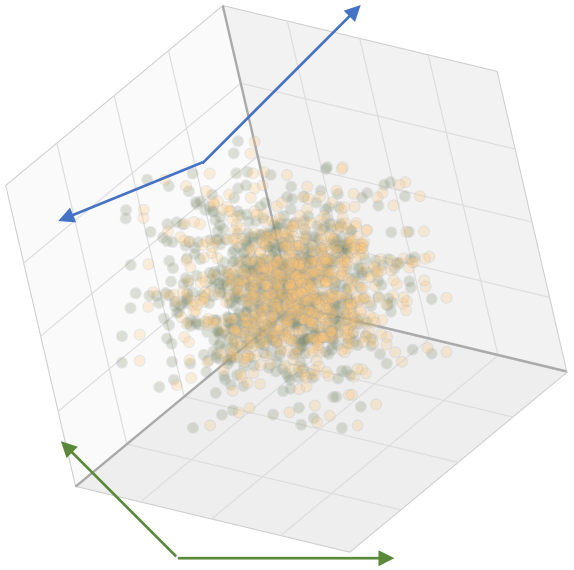

Semantic space

DISAMBIGUATION: IDENTIFYING MEANING AND DISTINGUISHING THE DIFFERENCES

A flexible framework able to:

- Reading the structure of a data sources and recognizing the concepts
- Identifying the definitions and distinguishing the differences
- Mapping them in a structure similar to our brain

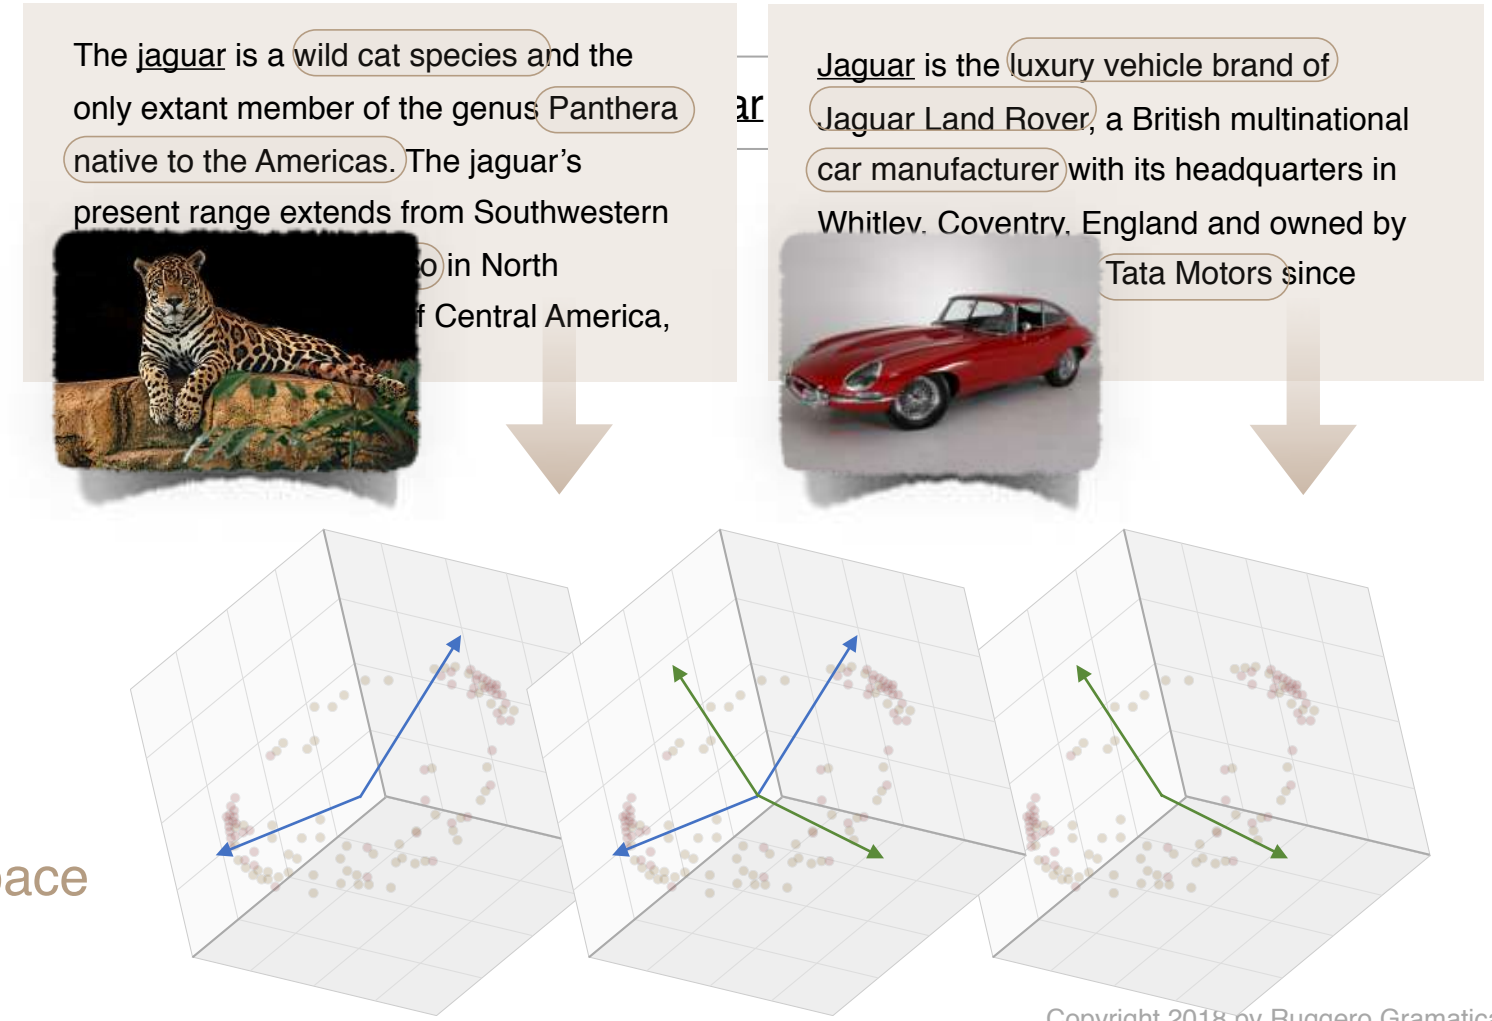

Semantic disambiguated space

THE CREATION OF A SEMANTIC SPACE

Iterating this process of projecting the Concept Universe leads to the construction of the

Semantic Space

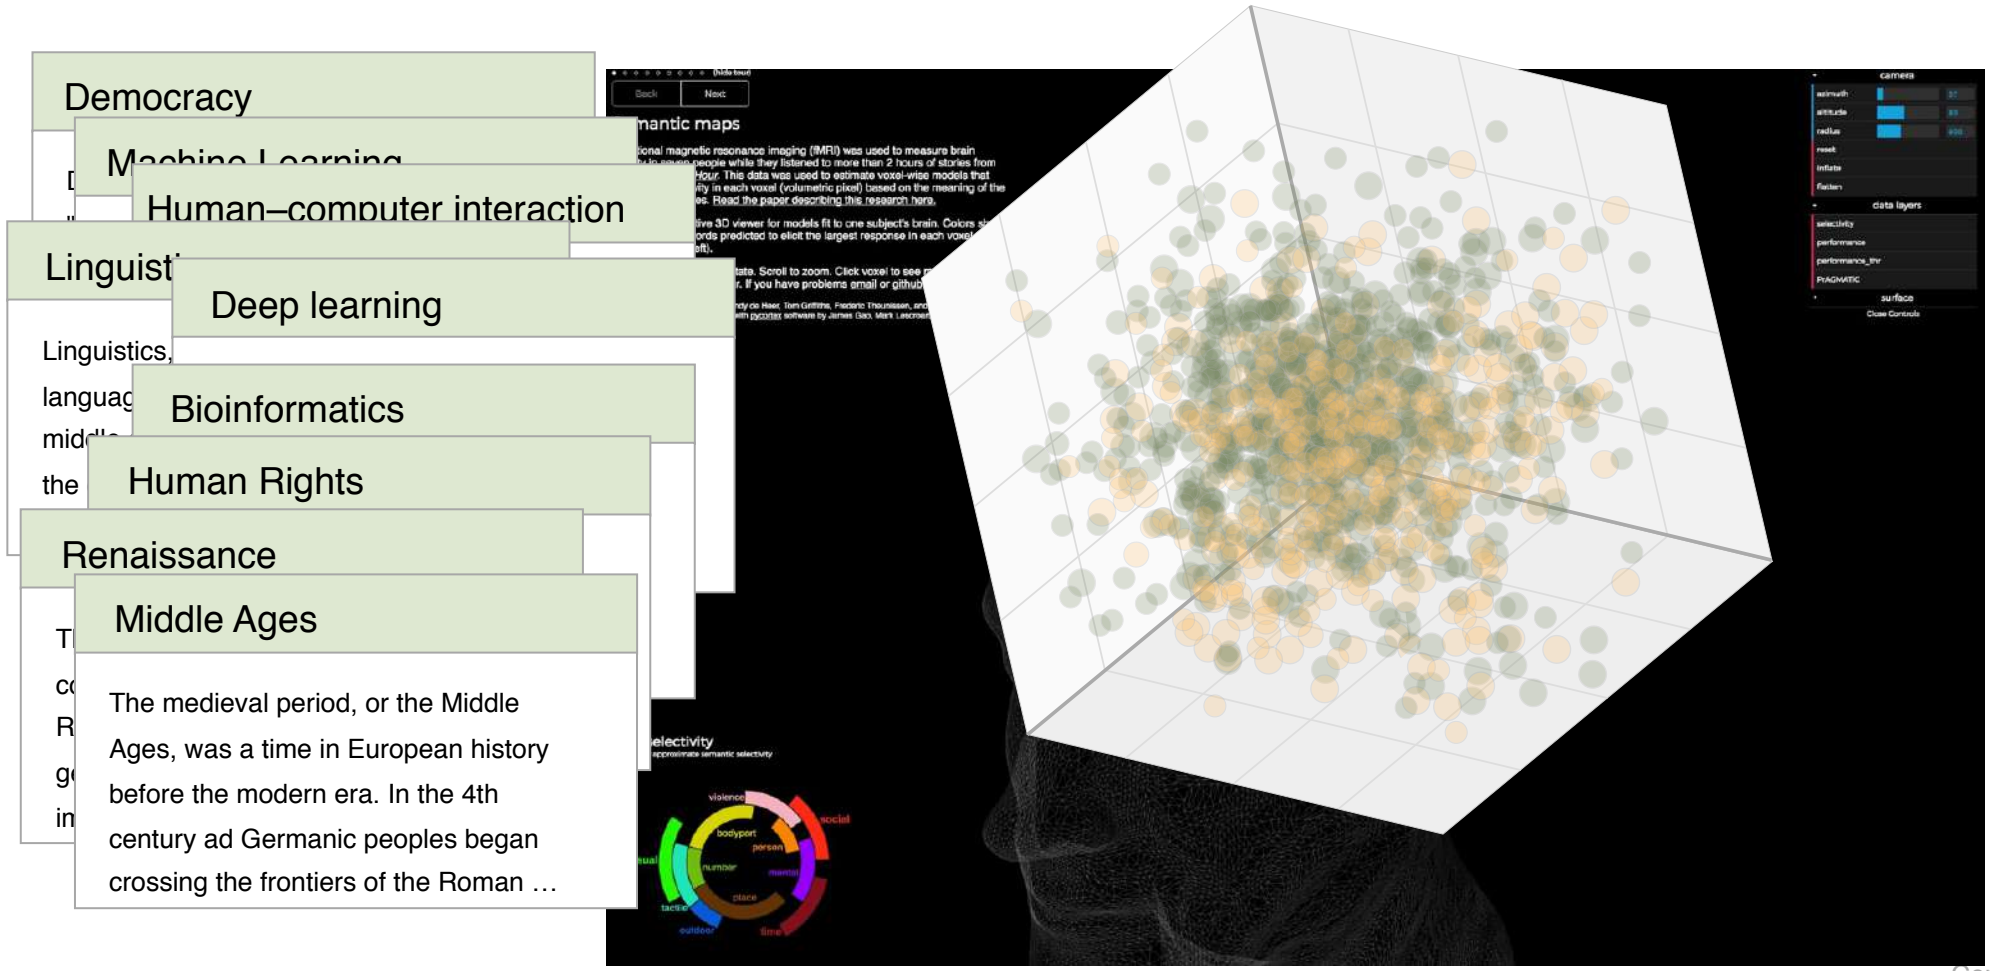

EVOLUTION OF CONCEPTS OVER TIME

Now we can map the evolution of the meaning of a concept.

How can we catch the evolution of concepts?

EXAMPLE

the concept “marriage"

200 years ago and today

Concepts vary over time.

- they arise
- they disappear
- they assume a different meaning

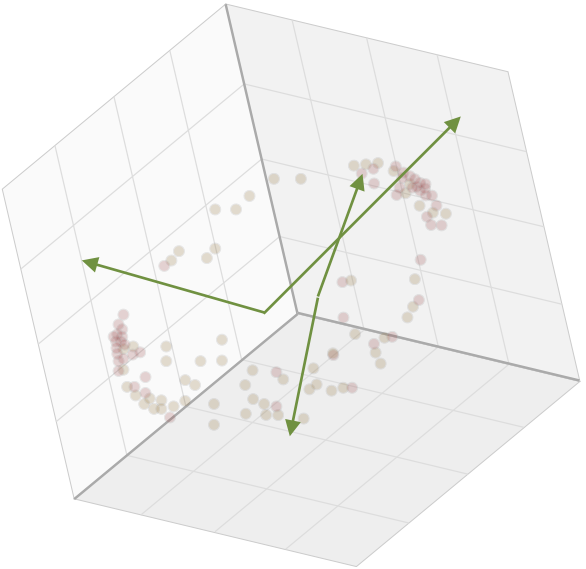

A language agnostic framework…..

MULTILINGUAL EMBEDDING: SEMANTIC SPACE ISOMORPHISM

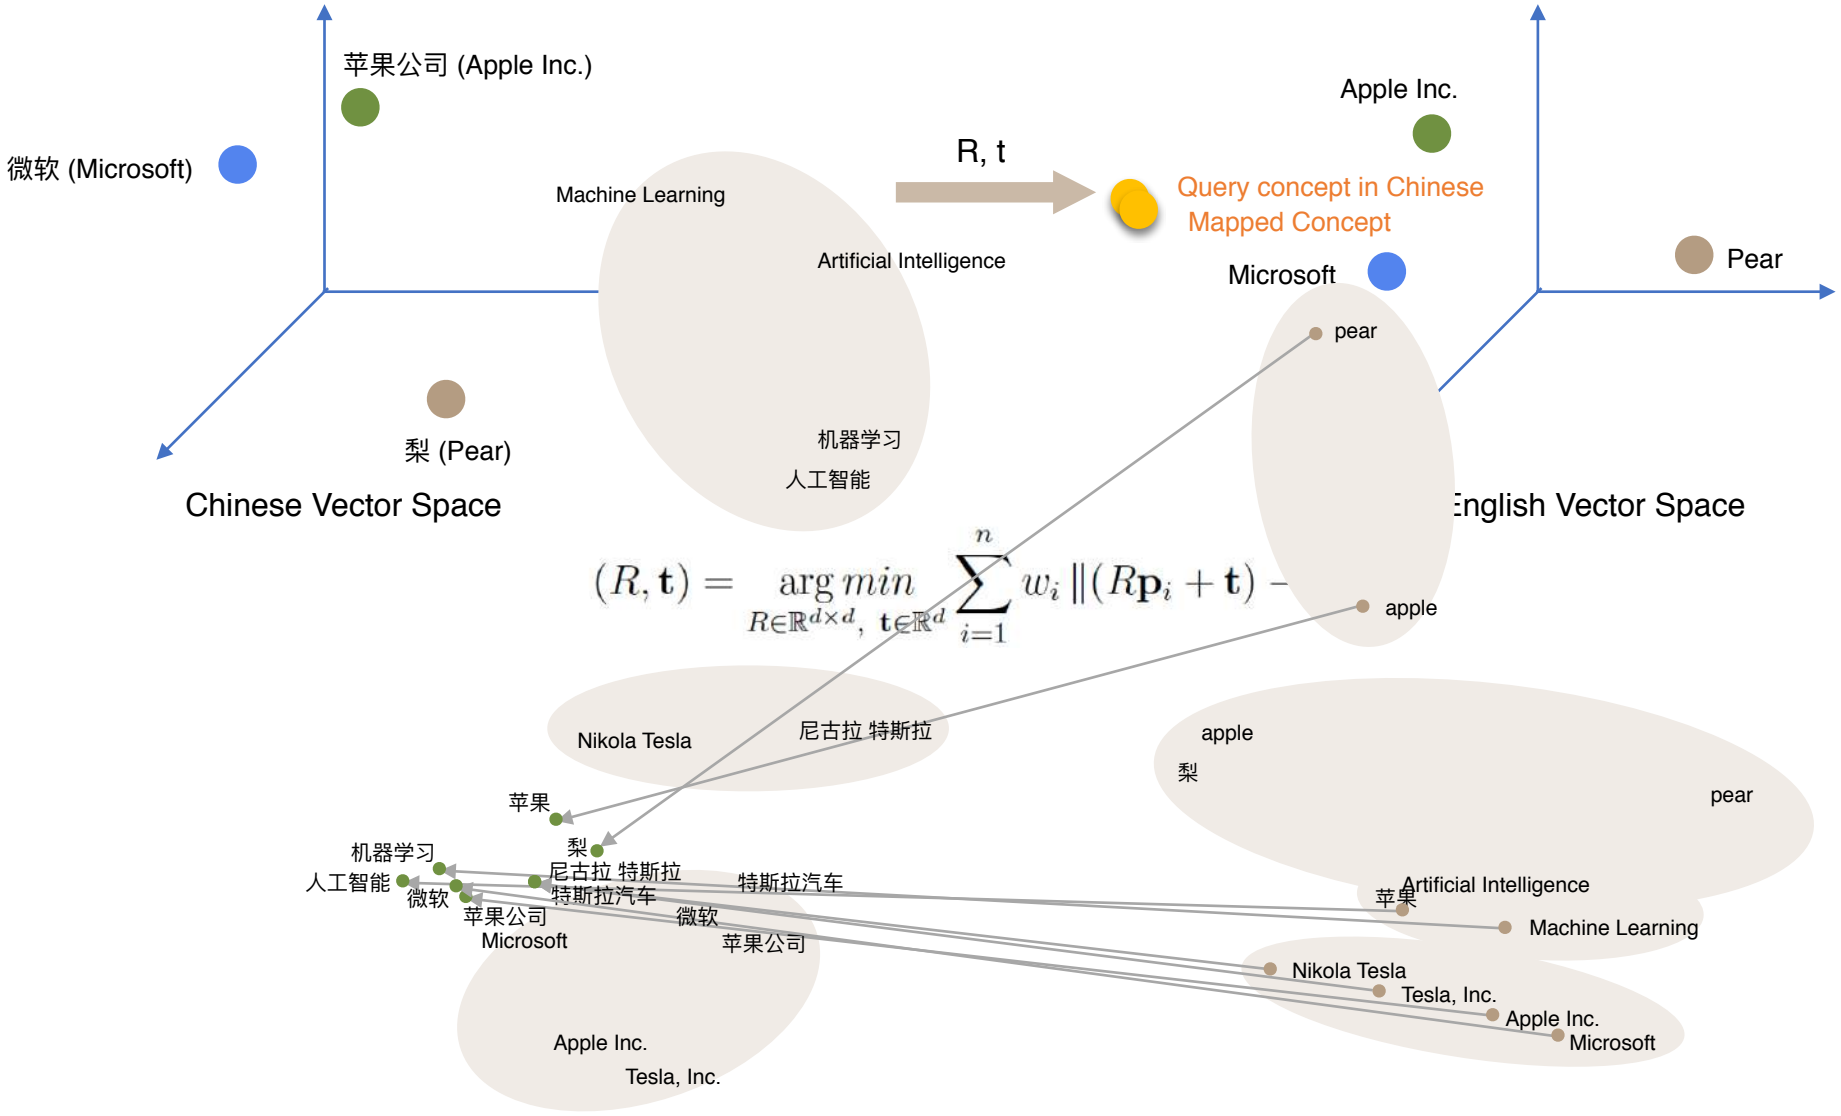

AN UNSUPERVISED INTELLIGENT FRAMEWORK EXTRACTING INFERENCES

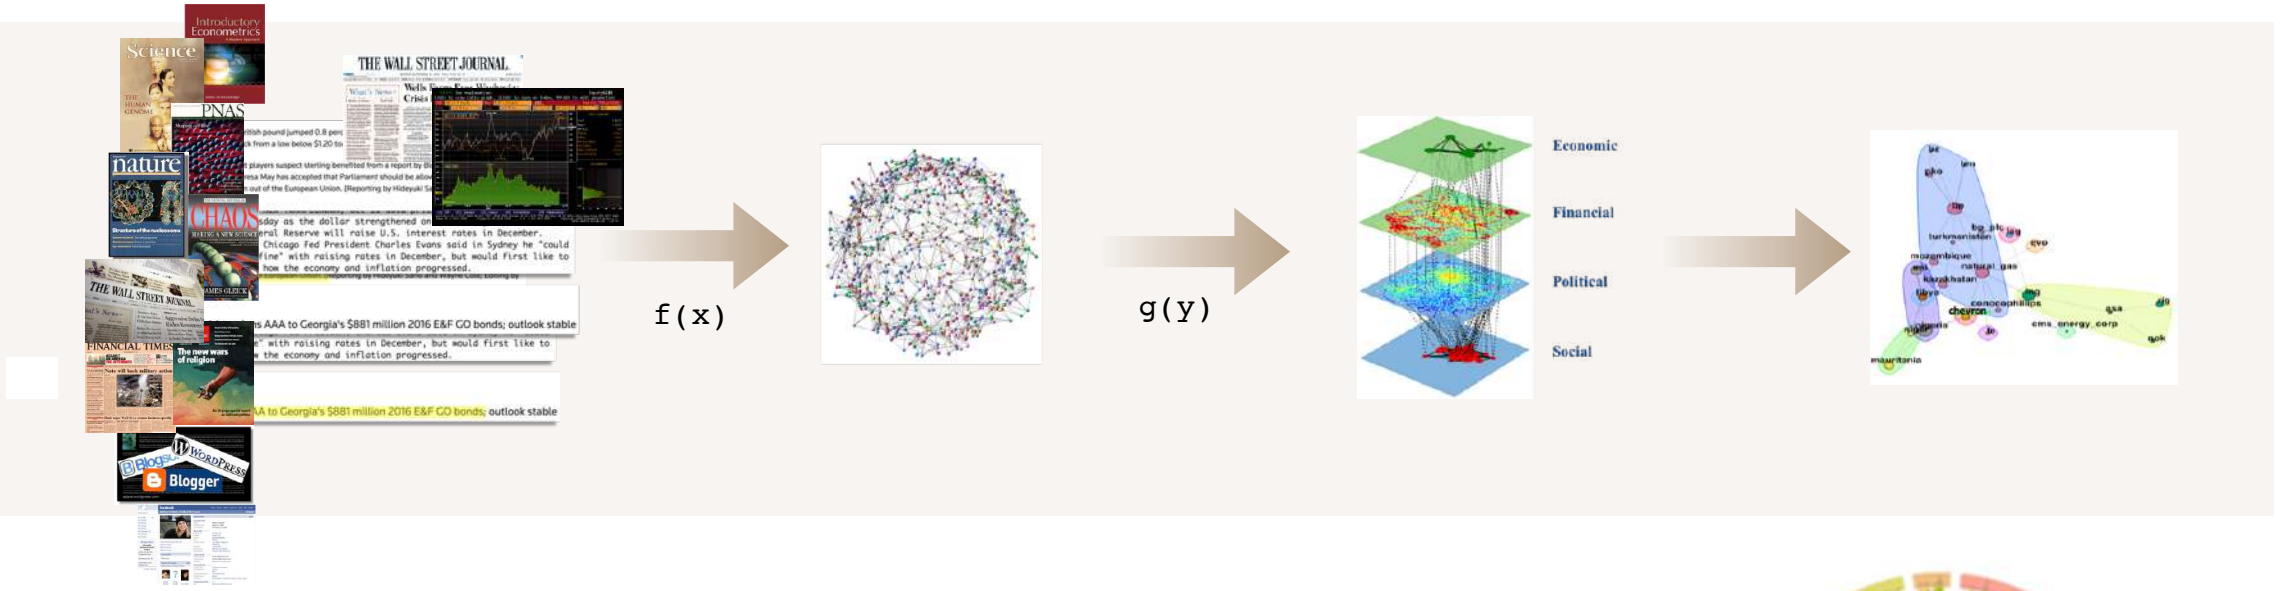

An intelligent framework that utilizes algorithms derived from Deep Learning neural networks allowing the extraction of both Quantitative and Semantic relationship in order to construct Inferential connections.

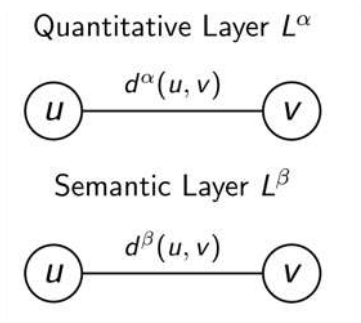

Inference Extraction

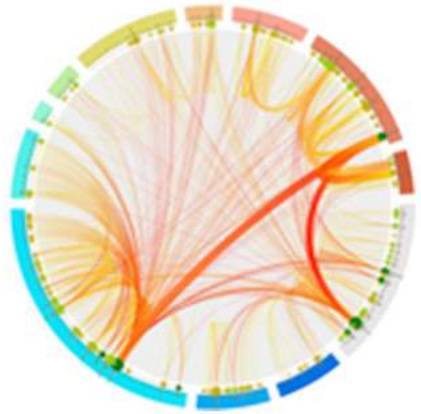

WHAT IS IT THEN A KNOWLEDGE GRAPH?

What it is not…!

- a static representation of entities
- a visualization tool
- a map of resources

What a knowledge graph is…

- A semantic projection of interconnected concepts
- a space where relationships are shown and explained
- A multi-domain framework where inferences are built

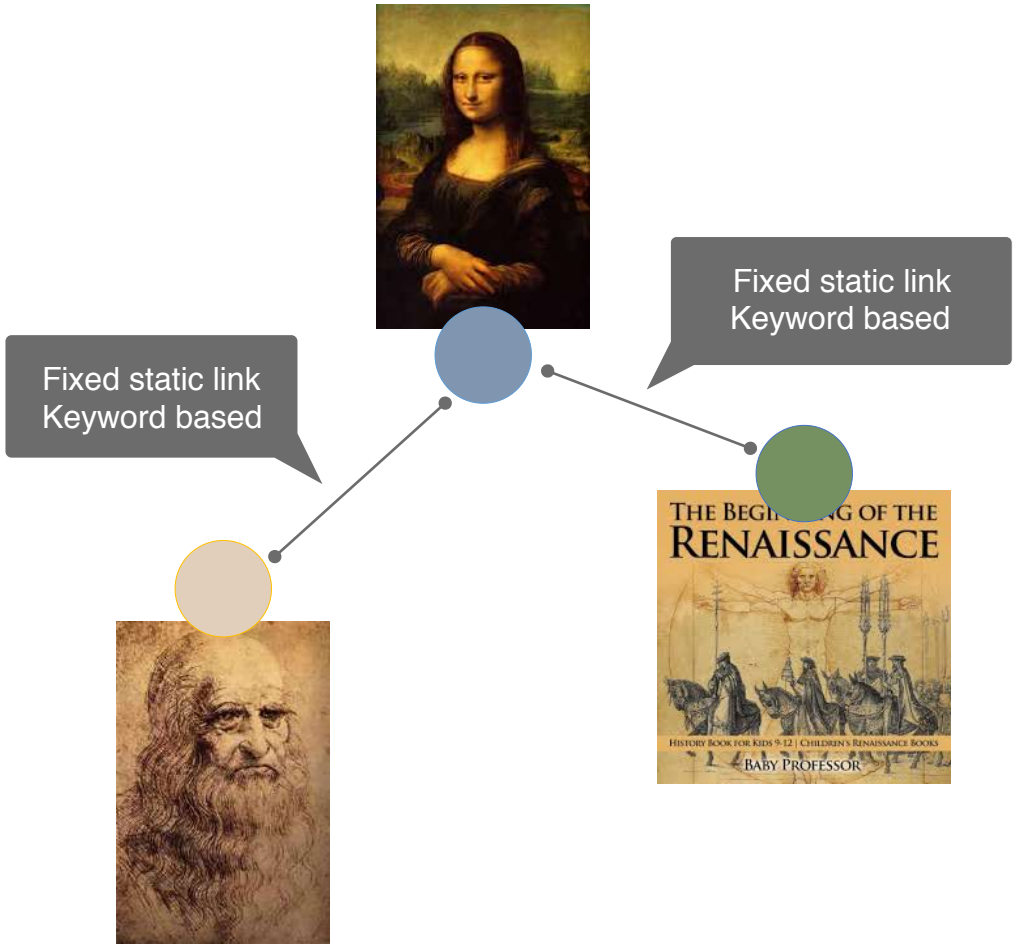

- STATIC REPRESENTATION OF INFORMATION…
- NOT MUCH KNOWLEDGE IN HERE…

WHAT IS IT THEN A KNOWLEDGE GRAPH?

A knowledge graph is a space where multi-dimensional inferences are built

- A semantic projection of interconnected concepts
- a space where relationships are shown and explained
- A multi-domain framework where inferences are built

Leonardo painted Mona Lisa in 1503 ca.

The Mona Lisa is a portrait of the wife of a Florentine merchant. The portrait was never delivered to its patron

An important copy of the Mona Lisa was recently discovered in the collection of the Prado in Madrid. …Scientific analysis revealed that the copy was likely painted by another a…ho sat beside Leonardo and

…

…rdo uses his characteristic sfumato to soften …s and create an atmospheric effect around the figure typical of the mature Renaissance style.

Leonardo combined Northern European innovations with Italian painting's understanding of the three dimensionality of the body and the perspectival treatment of the surrounding space.

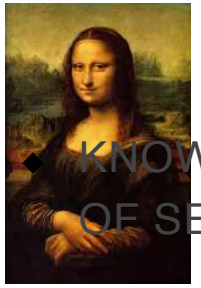

The *Mona Lisa* has become an icon of the *Renaissance art*.

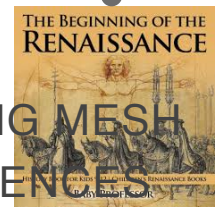

…

KNOWLEDGE GRAPH: AN EVER GROWING MESH OF SEMANTICALLY CORRELATED INFERENCES

AN INFINITE CONTENT UNIVERSE

INDUCTION OF A KNOWLEDGE GRAPH

DATA SOURCES

TECHNOLOGIES

Deep Neural Networks

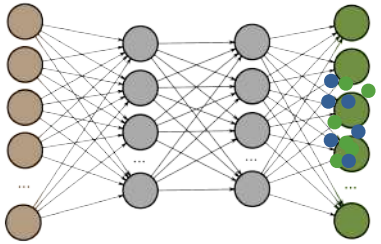

(Dynamic) topic models

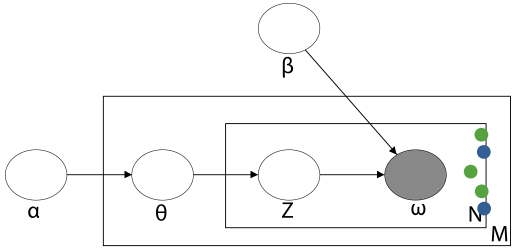

Stochastic learning

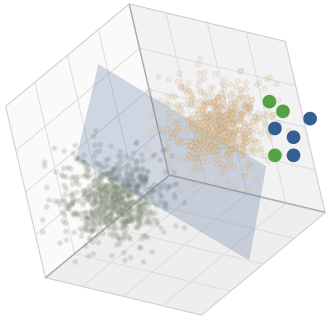

KNOWLEDGE GRAPH

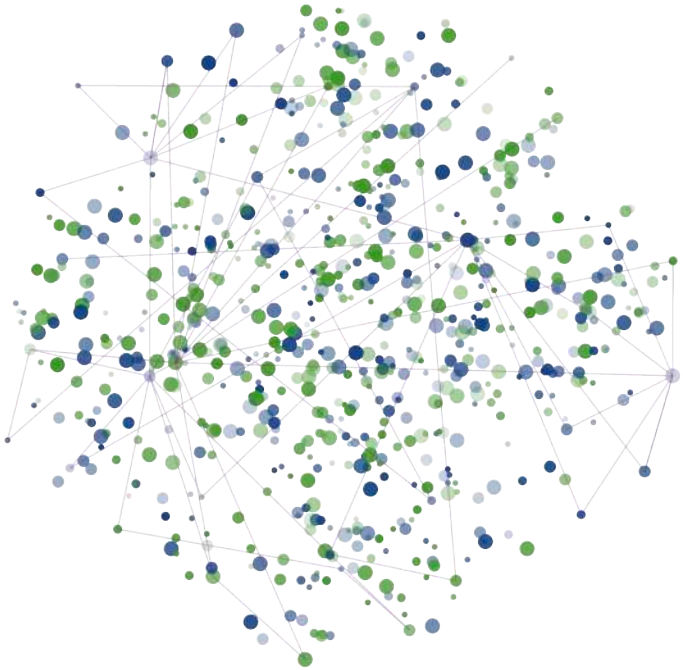

Knowledge Graph | POWERED BY Yewno

Yewno

black
vehicle
clothing
painting
son
social group
cartoon
animal
diagram
clip art
sea
transport
night
cademic conferenc
nature
web page
plant
residential area
house
food
structure

Transforming Information into Knowledge

A Dynamic and comprehensive Knowledge Graph,
Not just a theory…..
serving Knowledge Extraction

Yewno | Discover

And

Yewno | Unearth

Yewno | Discover

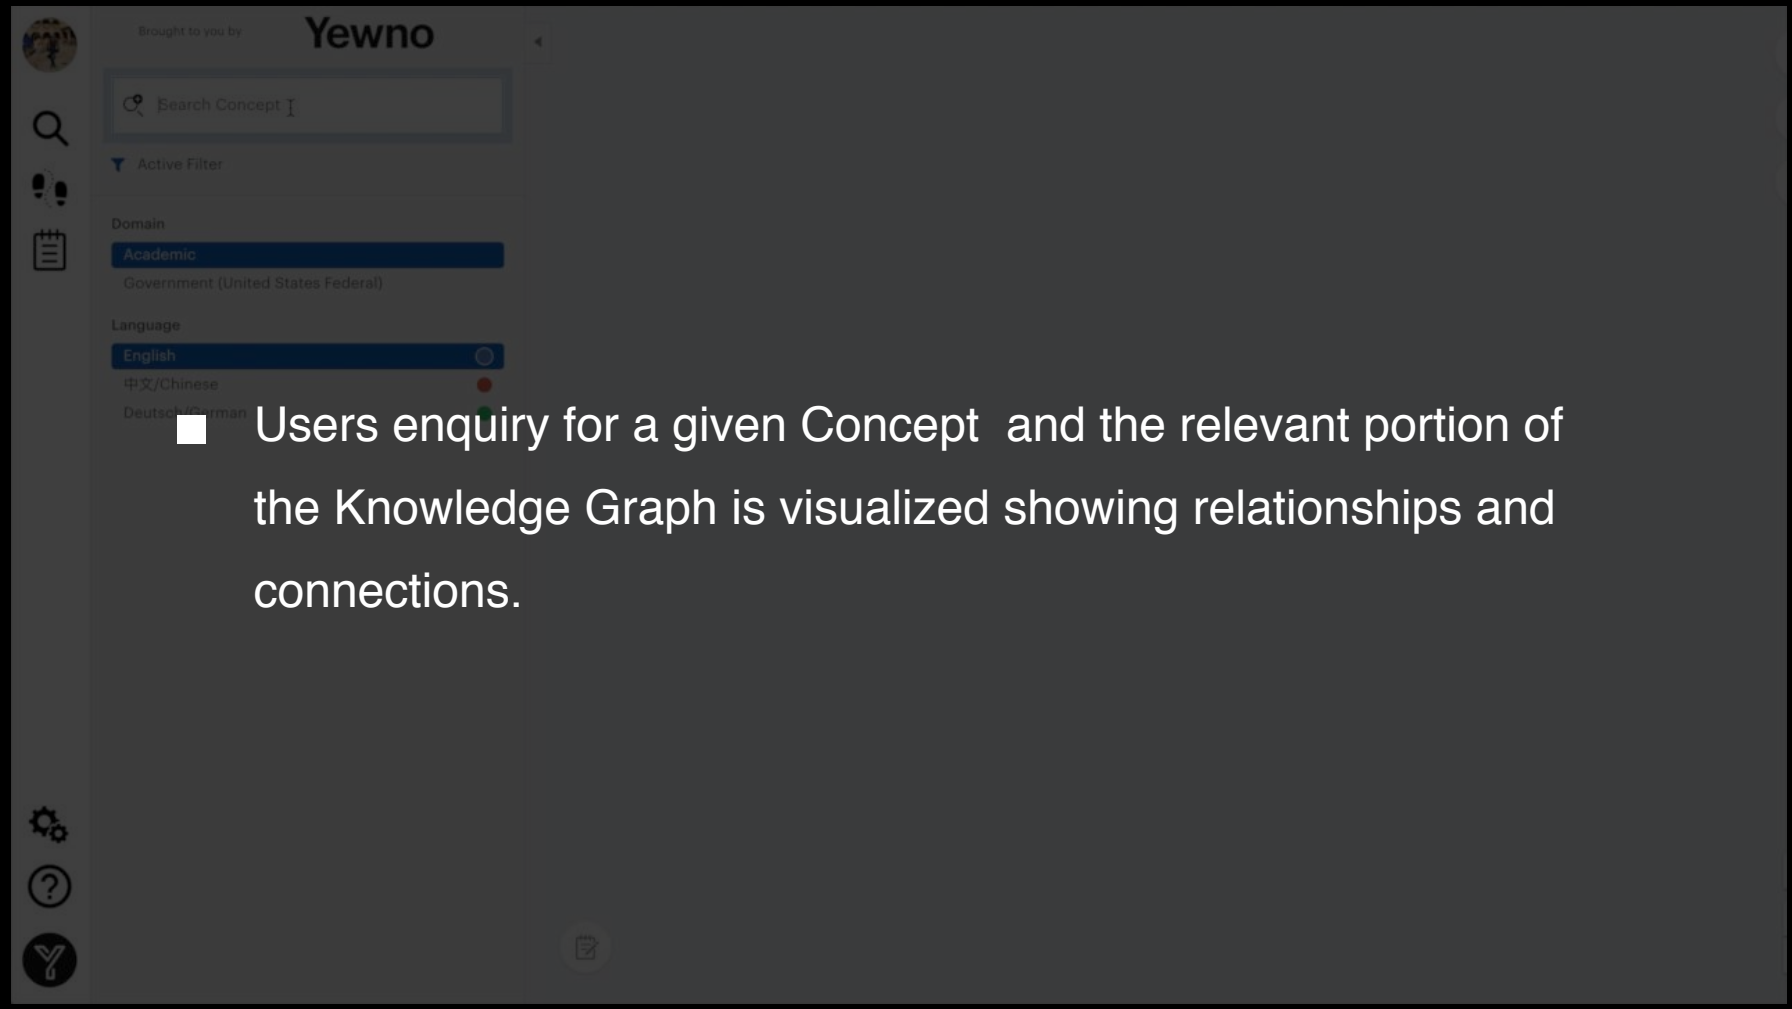

- Users enquiry for a given Concept and the relevant portion of the Knowledge Graph is visualized showing relationships and connections.

Yewno | Discover

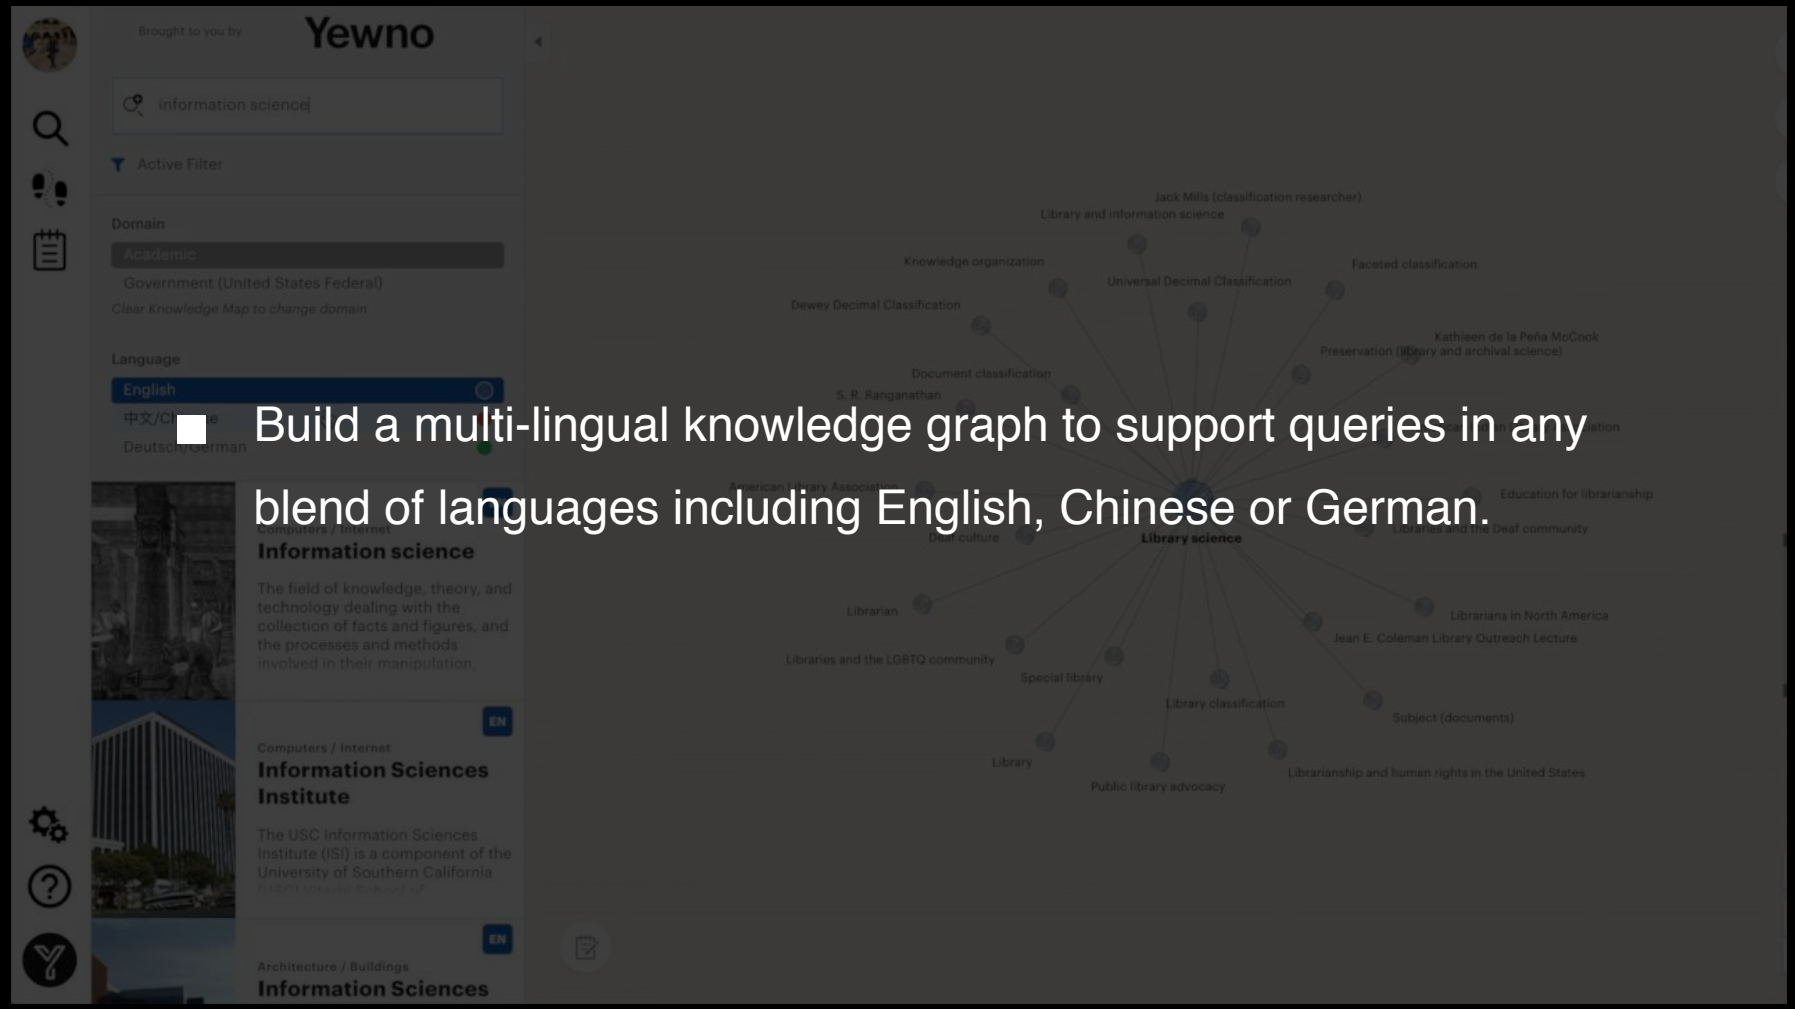

- Build a multi-lingual knowledge graph to support queries in any blend of languages including English, Chinese or German.

- Build a multi-lingual knowledge graph to support queries in any blend of languages including English, Chinese or German.

Yewno | Discover

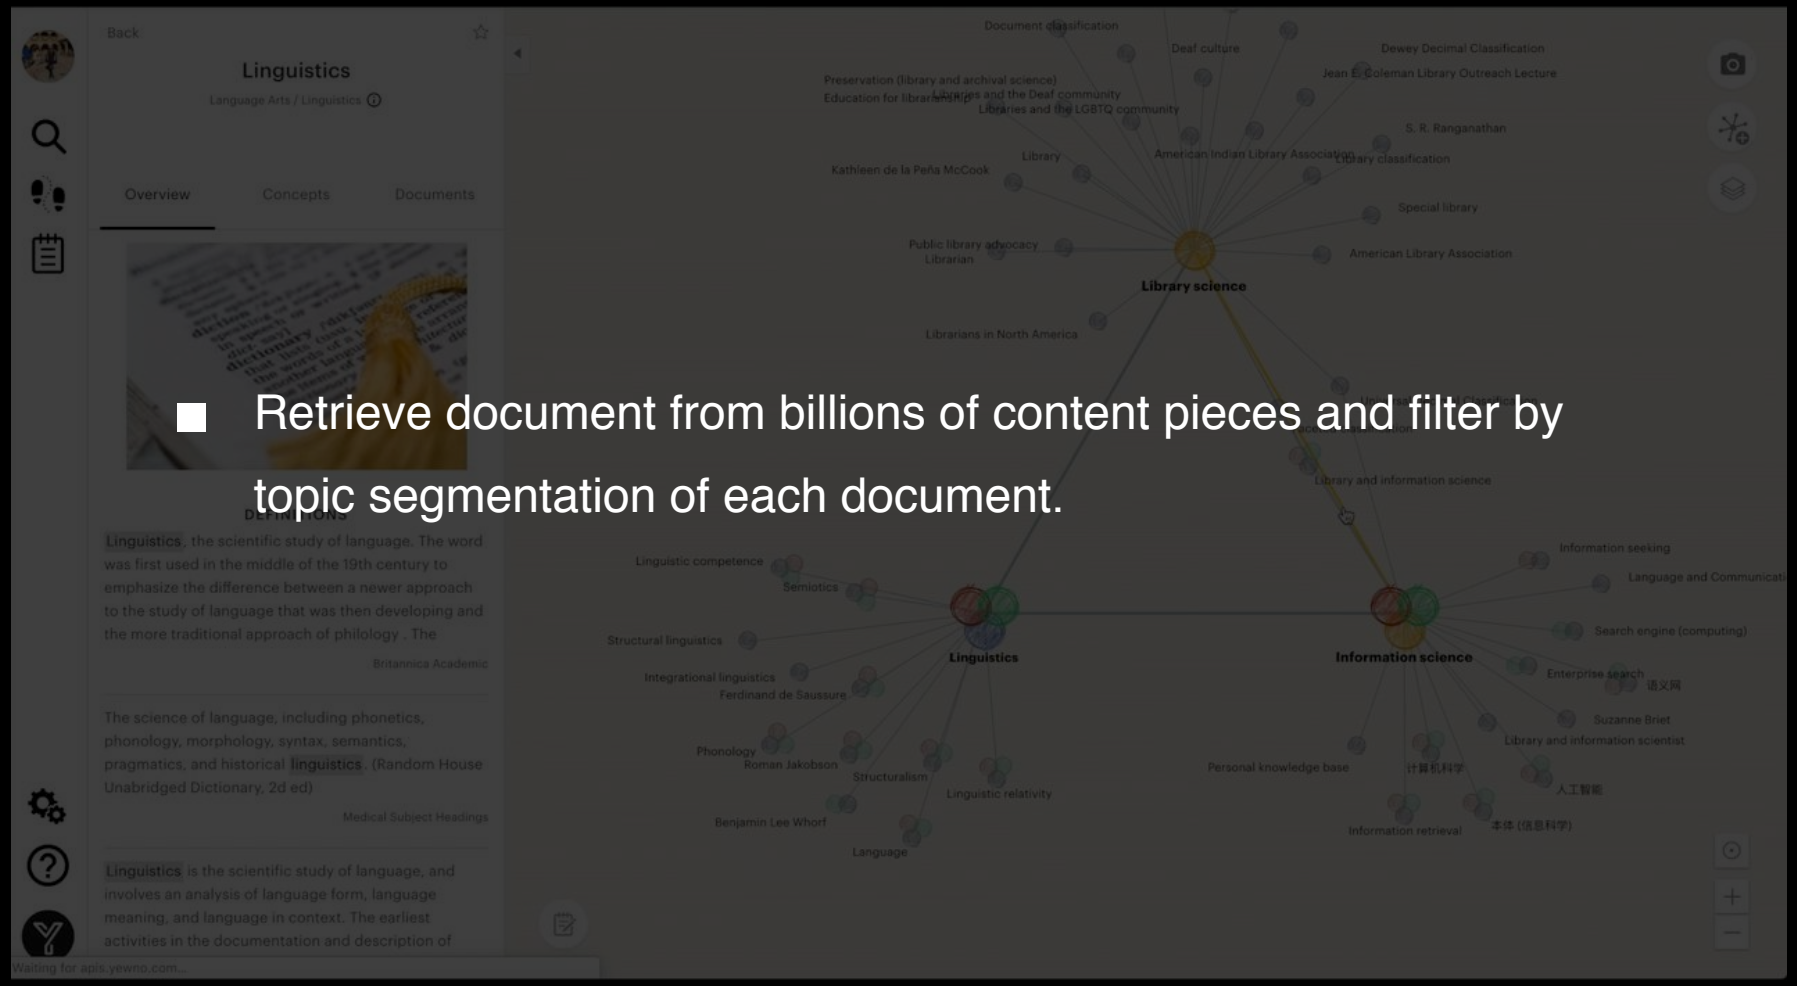

- Retrieve document from billions of content pieces and filter by topic segmentation of each document.

- Retrieve document from billions of content pieces and filter by topic segmentation of each document.

Yewno | Discover

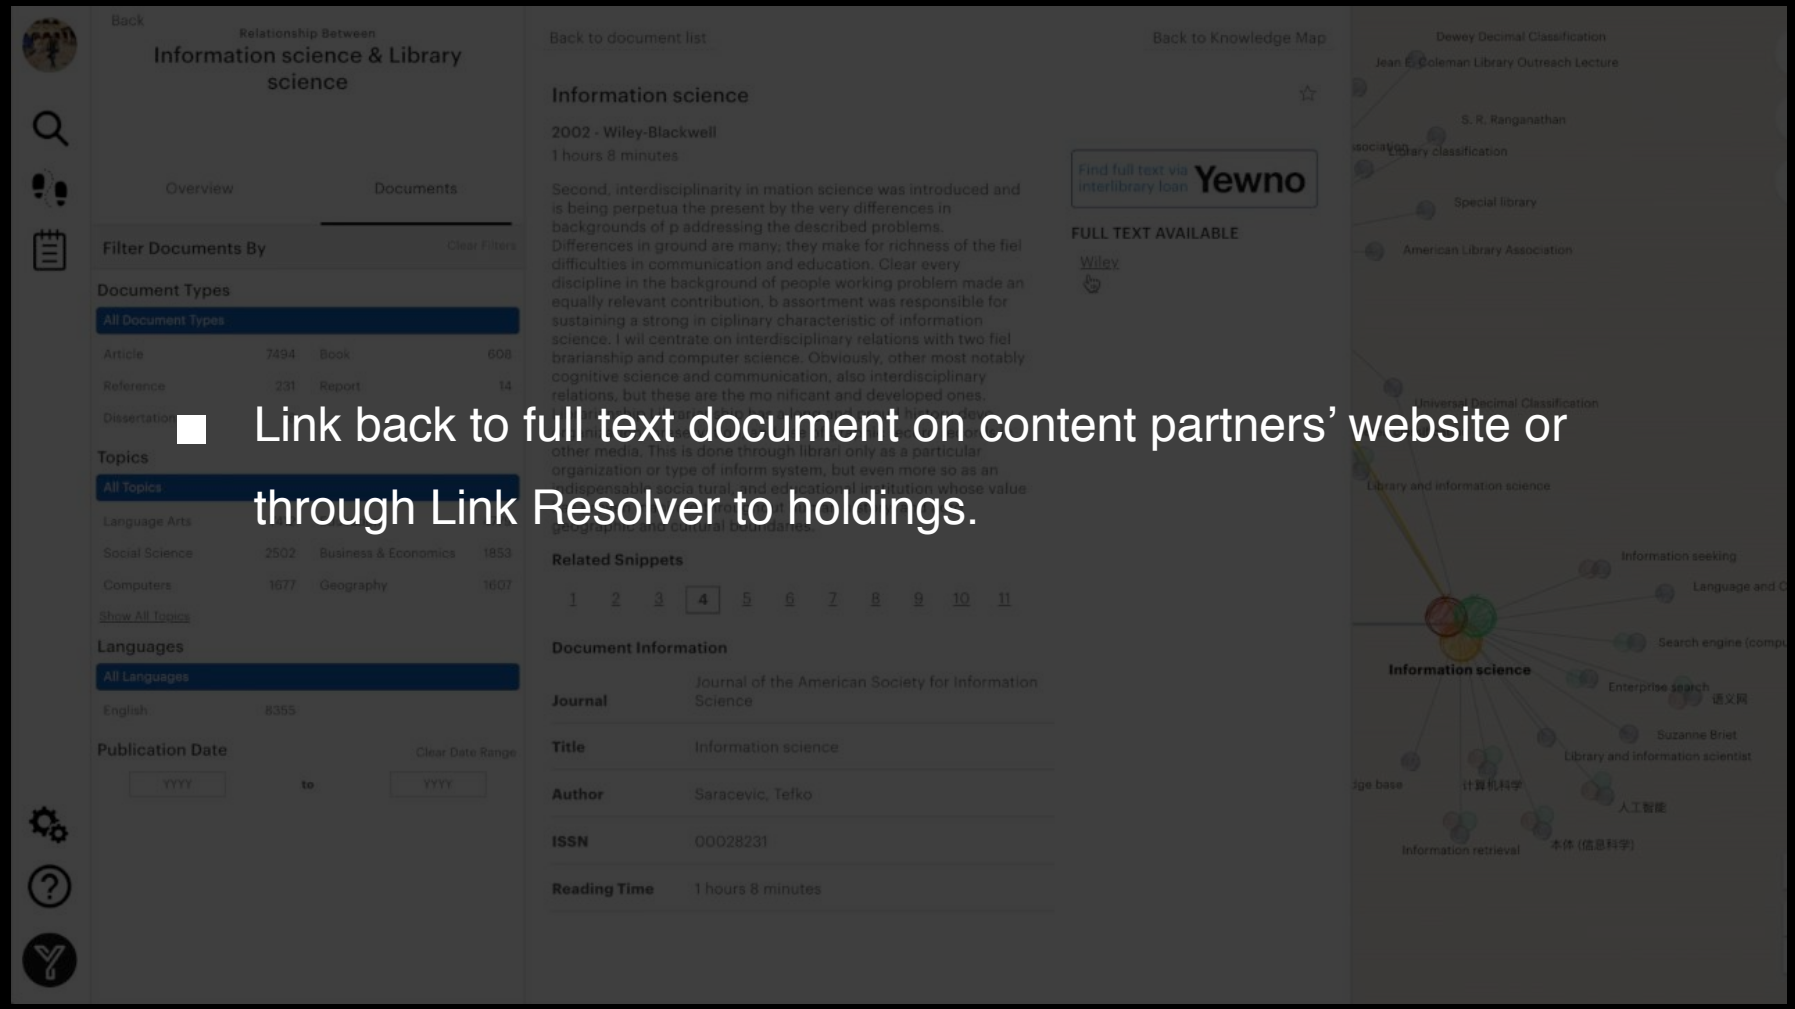

- Link back to full text document on content partners' website or through Link Resolver to holdings.

- Link back to full text document on content partners' website or through Link Resolver to holdings.

Yewno | Unearth

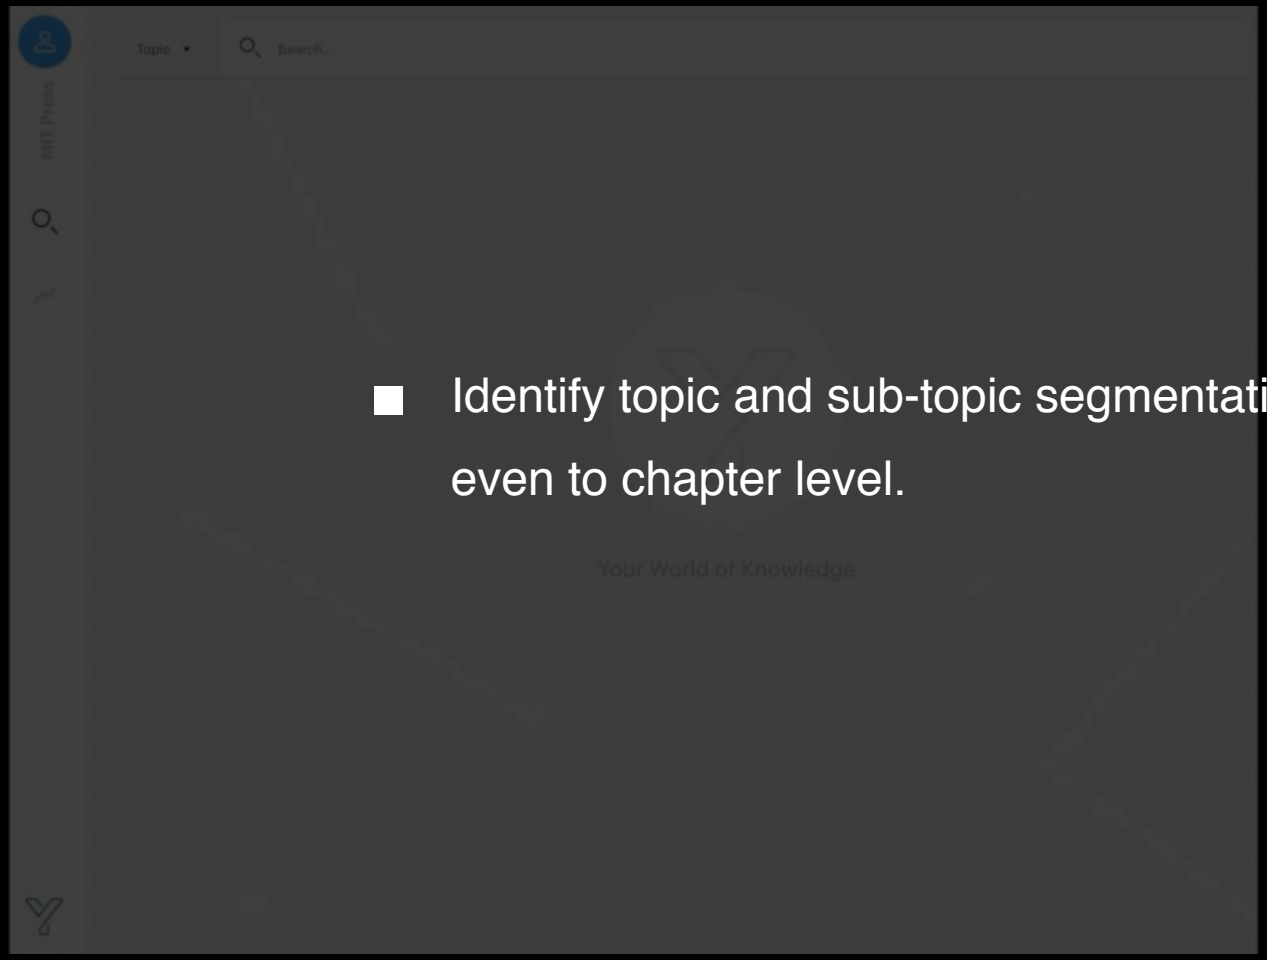

- Identify topic and sub-topic segmentation of each document even to chapter level.

- Identify topic and sub-topic segmentation of each document even to chapter level.

Yewno | Unearth

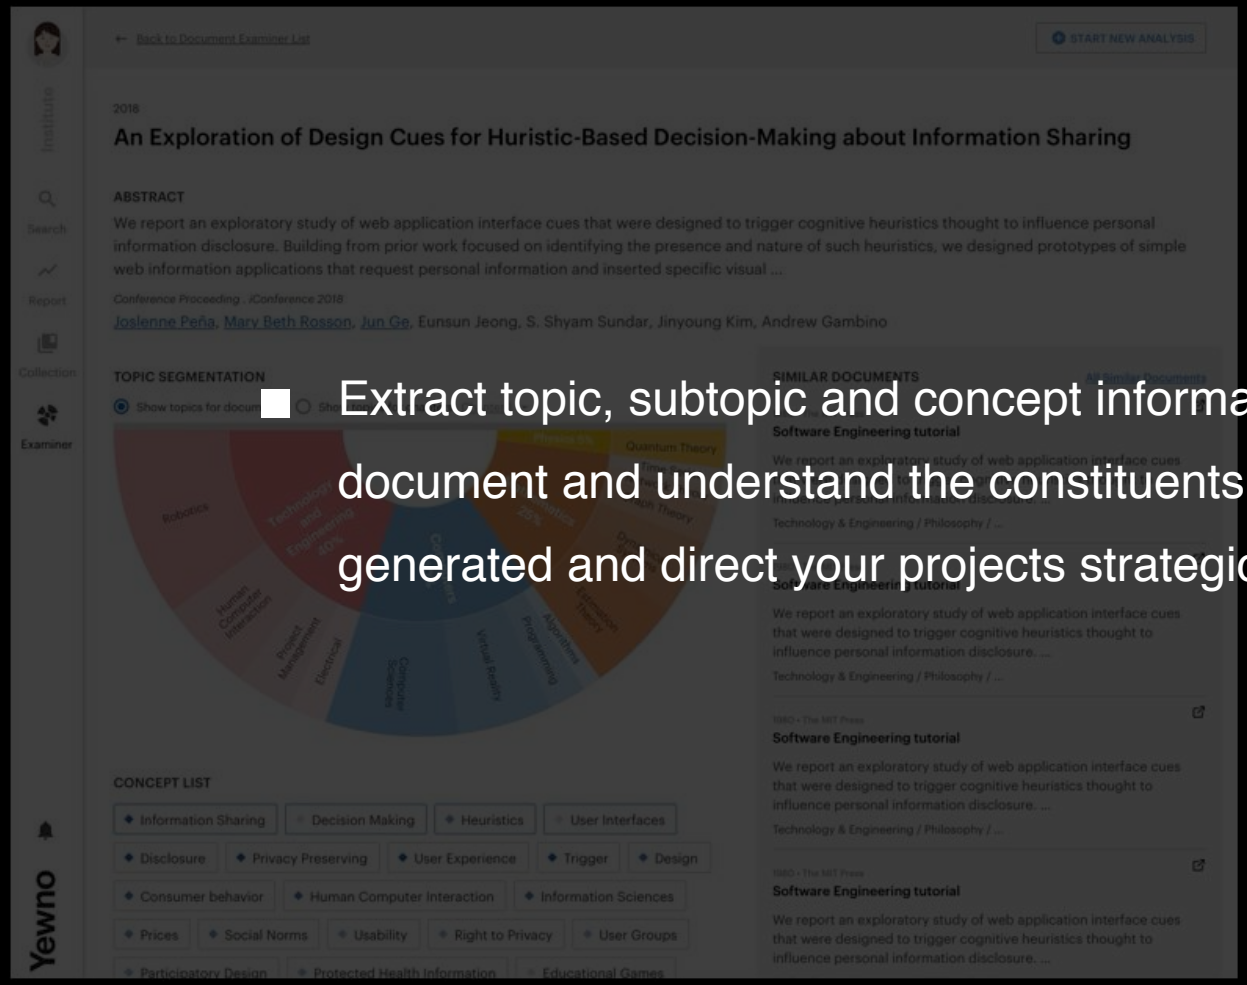

Yewno Knowledge Graph

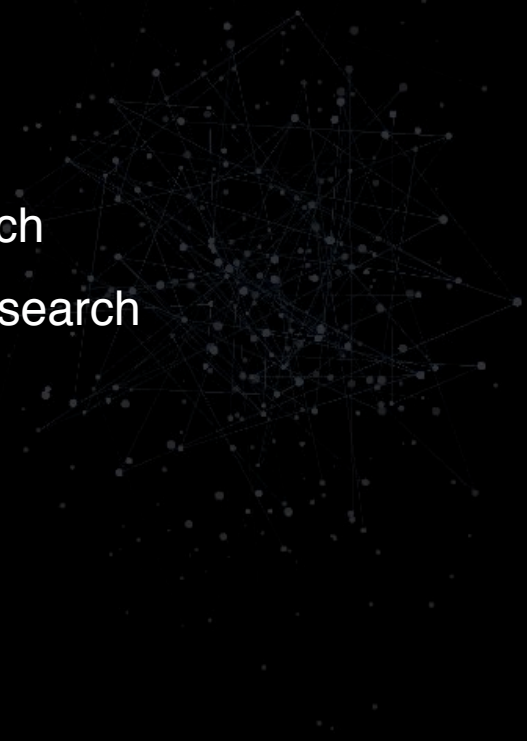

- Extract topic, subtopic and concept information of each document and understand the constituents of new research generated and direct your projects strategically.

- Extract topic, subtopic and concept information of each document and understand the constituents of new research generated and direct your projects strategically.

Yewno | Life Sciences

AN INSTRUMENT TO INVESTIGATE BIOMEDICAL RELATIONSHIPS

Knowledge Graph

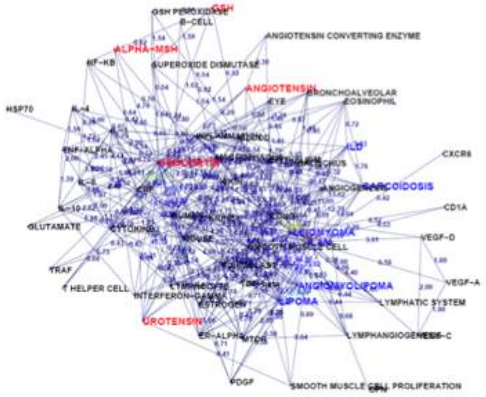

Derived Relationships

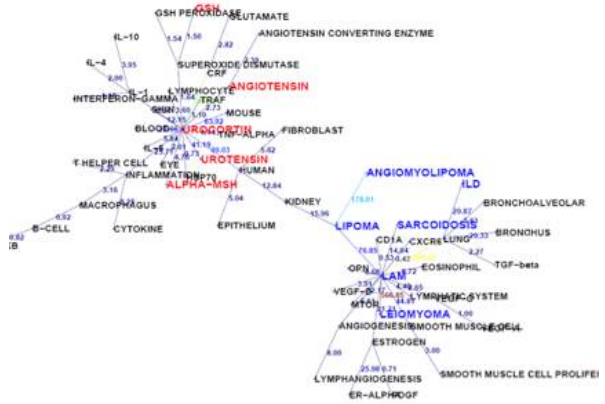

Interactive discovery applications

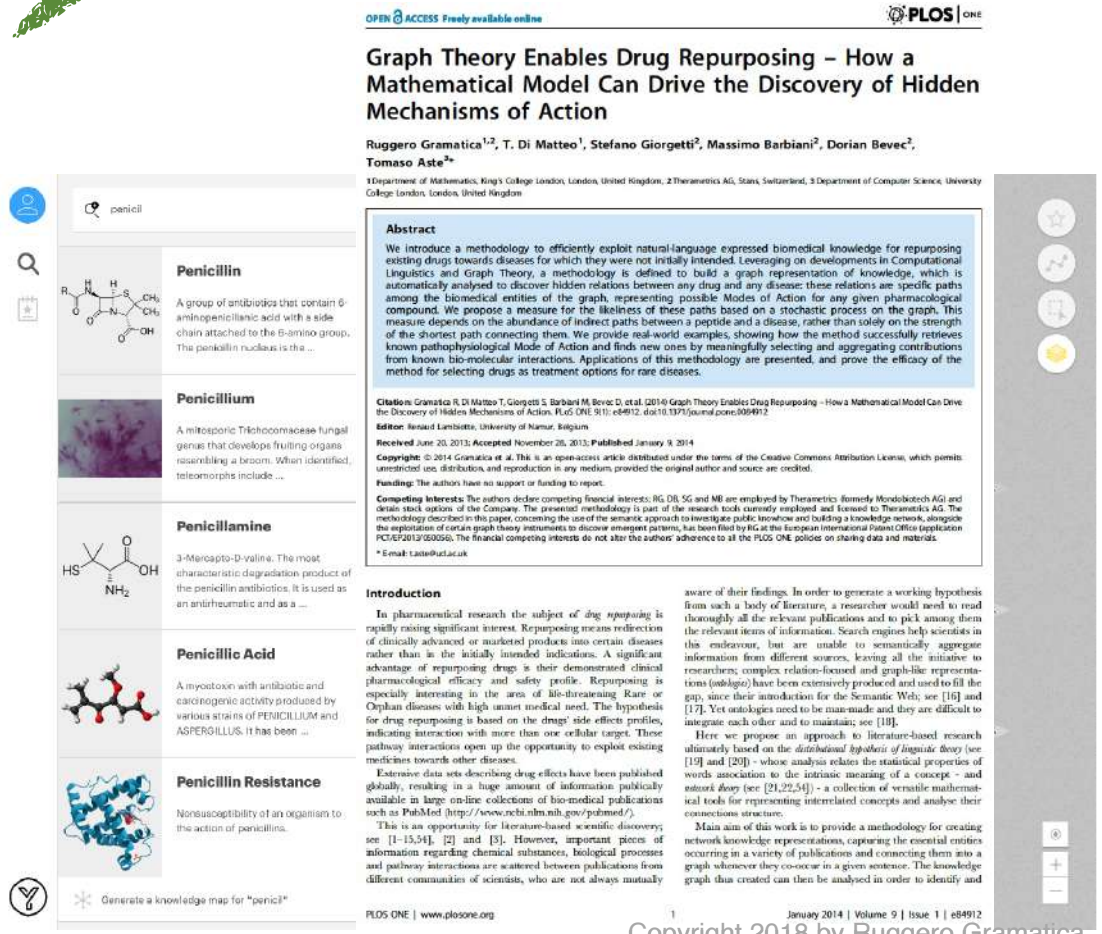

OPEN ACCESS Freely available online

PLOS ONE

Graph Theory Enables Drug Repurposing – How a Mathematical Model Can Drive the Discovery of Hidden Mechanisms of Action

Ruggero Gramatica[1,2], T. Di Matteo[1], Stefano Giorgetti[2], Massimo Barbiani[2], Dorian Bevec[2], Tomaso Aste[3]*

[1] Department of Mathematics, King's College London, London, United Kingdom, [2] Thermetrics AG, Stans, Switzerland, [3] Department of Computer Science, University College London, London, United Kingdom

Abstract

We introduce a methodology to efficiently exploit natural-language expressed biomedical knowledge for repurposing existing drugs towards diseases for which they were not initially intended. Leveraging on developments in Computational Linguistics and Graph Theory, a methodology is defined to build a graph representation of knowledge, which is automatically analysed to discover hidden relations between any drug and any disease: these relations are specific paths among the biomedical entities of the graph, representing possible Modes of Action for any given pharmacological compound. We propose a measure for the likeliness of these paths based on a stochastic process on the graph. This measure depends on the abundance of indirect paths between a peptide and a disease, rather than solely on the strength of the shortest path connecting them. We provide real-world examples, showing how the method successfully retrieves known pathophysiological Mode of Action and finds new ones by meaningfully selecting and aggregating contributions from known bio-molecular interactions. Applications of this methodology are presented, and prove the efficacy of the method for selecting drugs as treatment options for rare diseases.

Citation: Gramatica R, Di Matteo T, Giorgetti S, Barbiani M, Bevec D, et al. (2014) Graph Theory Enables Drug Repurposing – How a Mathematical Model Can Drive the Discovery of Hidden Mechanisms of Action. PLoS ONE 9(1): e84912. doi:10.1371/journal.pone.0084912

Editor: Renaud Lambiotte, University of Namur, Belgium

Received June 20, 2013; **Accepted** November 26, 2013; **Published** January 9, 2014

Funding: The authors have no support or funding to report.

Competing Interests: The authors declare competing financial interests: RG, DB, SG and MB are employed by Thermetrics (formerly Mondobiotech AG) and detain stock options of the Company. The presented methodology is part of the research tools currently employed and licensed to Thermetrics AG. The methodology described in this paper, concerning the use of the semantic approach to investigate public knowhow and building a knowledge network, alongside the exploitation of certain graph theory instruments to discover emergent patterns, has been filed by RG at the European International Patent Office (application PCT/EP2013/050056). The financial competing interests do not alter the authors' adherence to all the PLOS ONE policies on sharing data and materials.

* E-mail: t.aste@ucl.ac.uk

Introduction

In pharmaceutical research the subject of *drug repurposing* is rapidly raising significant interest. Repurposing means redirection of clinically advanced or marketed products into certain diseases rather than in the initially intended indications. A significant advantage of repurposing drugs is their demonstrated clinical pharmacological efficacy and safety profile. Repurposing is especially interesting in the area of life-threatening Rare or Orphan diseases with high unmet medical need. The hypothesis for drug repurposing is based on the drugs' side effects profiles, indicating interaction with more than one cellular target. These pathway interactions open up the opportunity to exploit existing medicines towards other diseases.

Extensive data sets describing drug effects have been published globally, resulting in a huge amount of information publically available in large on-line collections of bio-medical publications such as PubMed (http://www.ncbi.nlm.nih.gov/pubmed/).

This is an opportunity for literature-based scientific discovery; see [1–15,54], [2] and [3]. However, important pieces of information regarding chemical substances, biological processes and pathway interactions are scattered between publications from different communities of scientists, who are not always mutually aware of their findings. In order to generate a working hypothesis from such a body of literature, a researcher would need to read thoroughly all the relevant publications and to pick among them the relevant items of information. Search engines help scientists in this endeavour, but are unable to semantically aggregate information from different sources, leaving all the initiative to researchers; complex relation-focused and graph-like representations (*ontologies*) have been extensively produced and used to fill the gap, since their introduction for the Semantic Web; see [16] and [17]. Yet ontologies need to be man-made and they are difficult to integrate each other and to maintain; see [18].

Here we propose an approach to literature-based research ultimately based on the *distributional hypothesis of linguistic theory* (see [19] and [20]) - whose analysis relates the statistical properties of words association to the intrinsic meaning of a concept - and *network theory* (see [21,22,54]) - a collection of versatile mathematical tools for representing interrelated concepts and analyse their connections structure.

Main aim of this work is to provide a methodology for creating network knowledge representations, capturing the essential entities occurring in a variety of publications and connecting them into a graph whenever they co-occur in a given sentence. The knowledge graph thus created can then be analysed in order to identify and

(*) *Graph theory enables drug repurposing – How a mathematical model can drive the discovery of hidden Mechanisms of Action* (Ruggero Gramatica et alii 2014 - PLOS|ONE)

Yewno

Yewno | Life Science

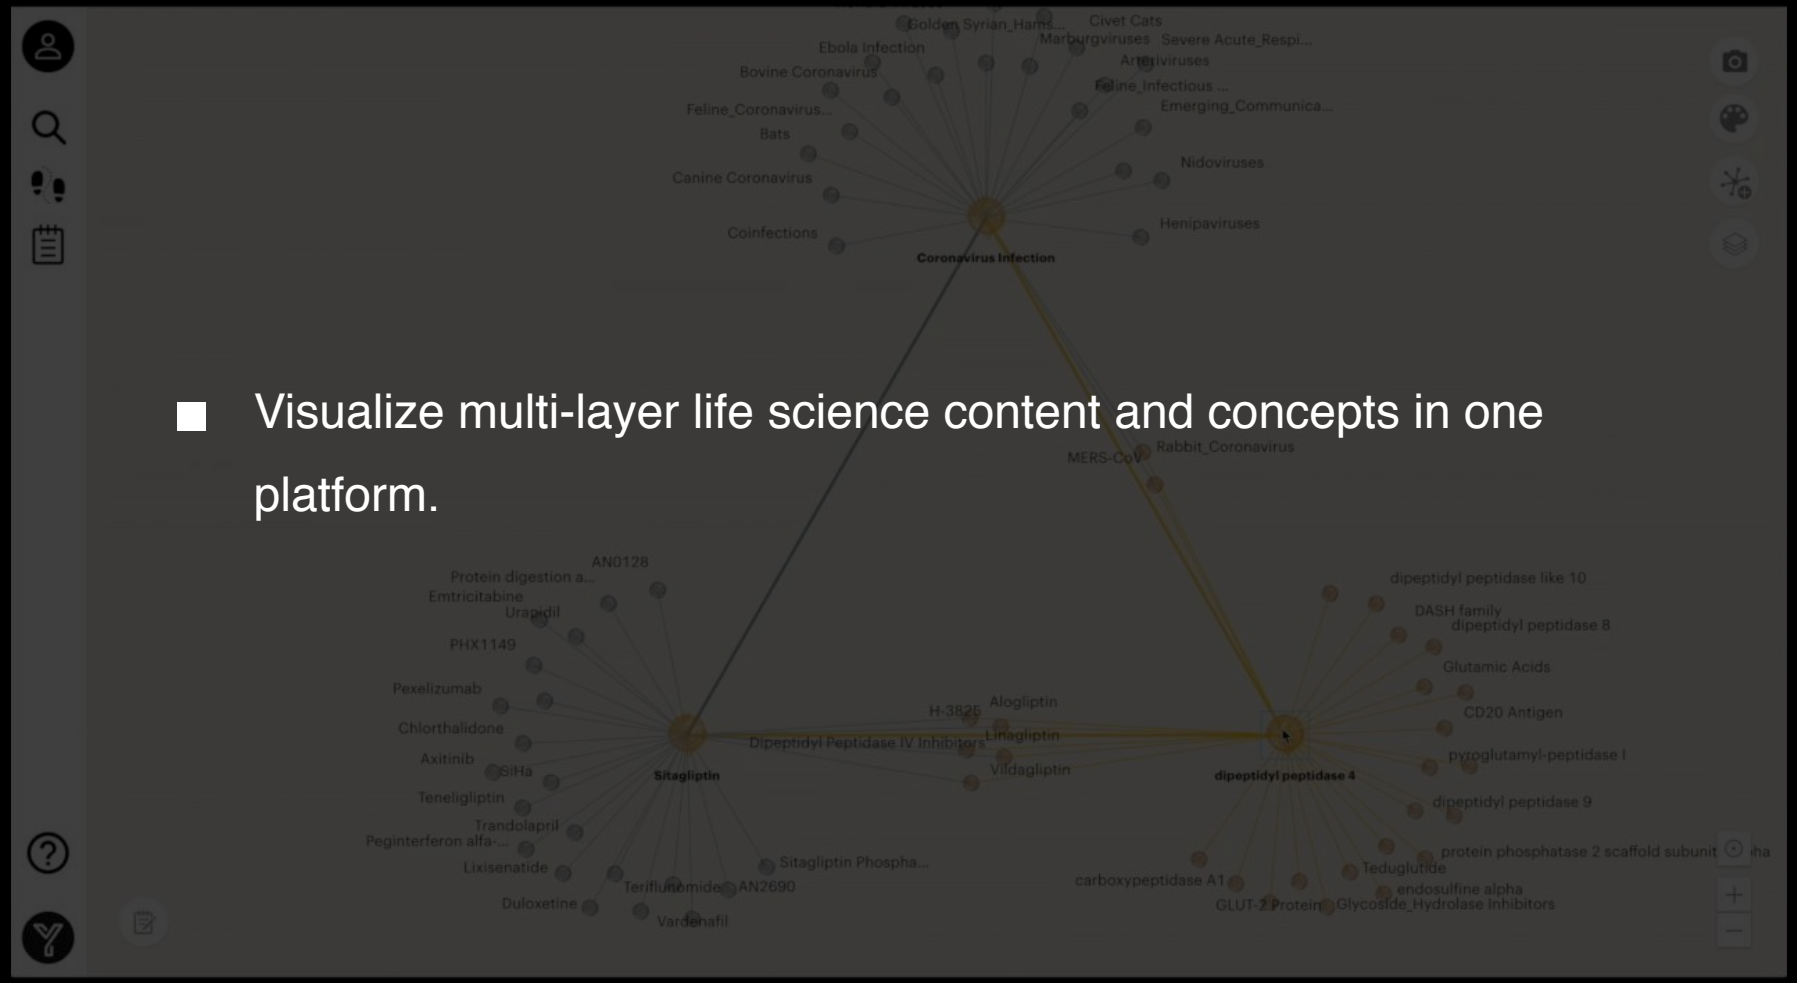

- Visualize multi-layer life science content and concepts in one platform.

- Visualize multi-layer life science content and concepts in one platform.

Yewno | Life Science

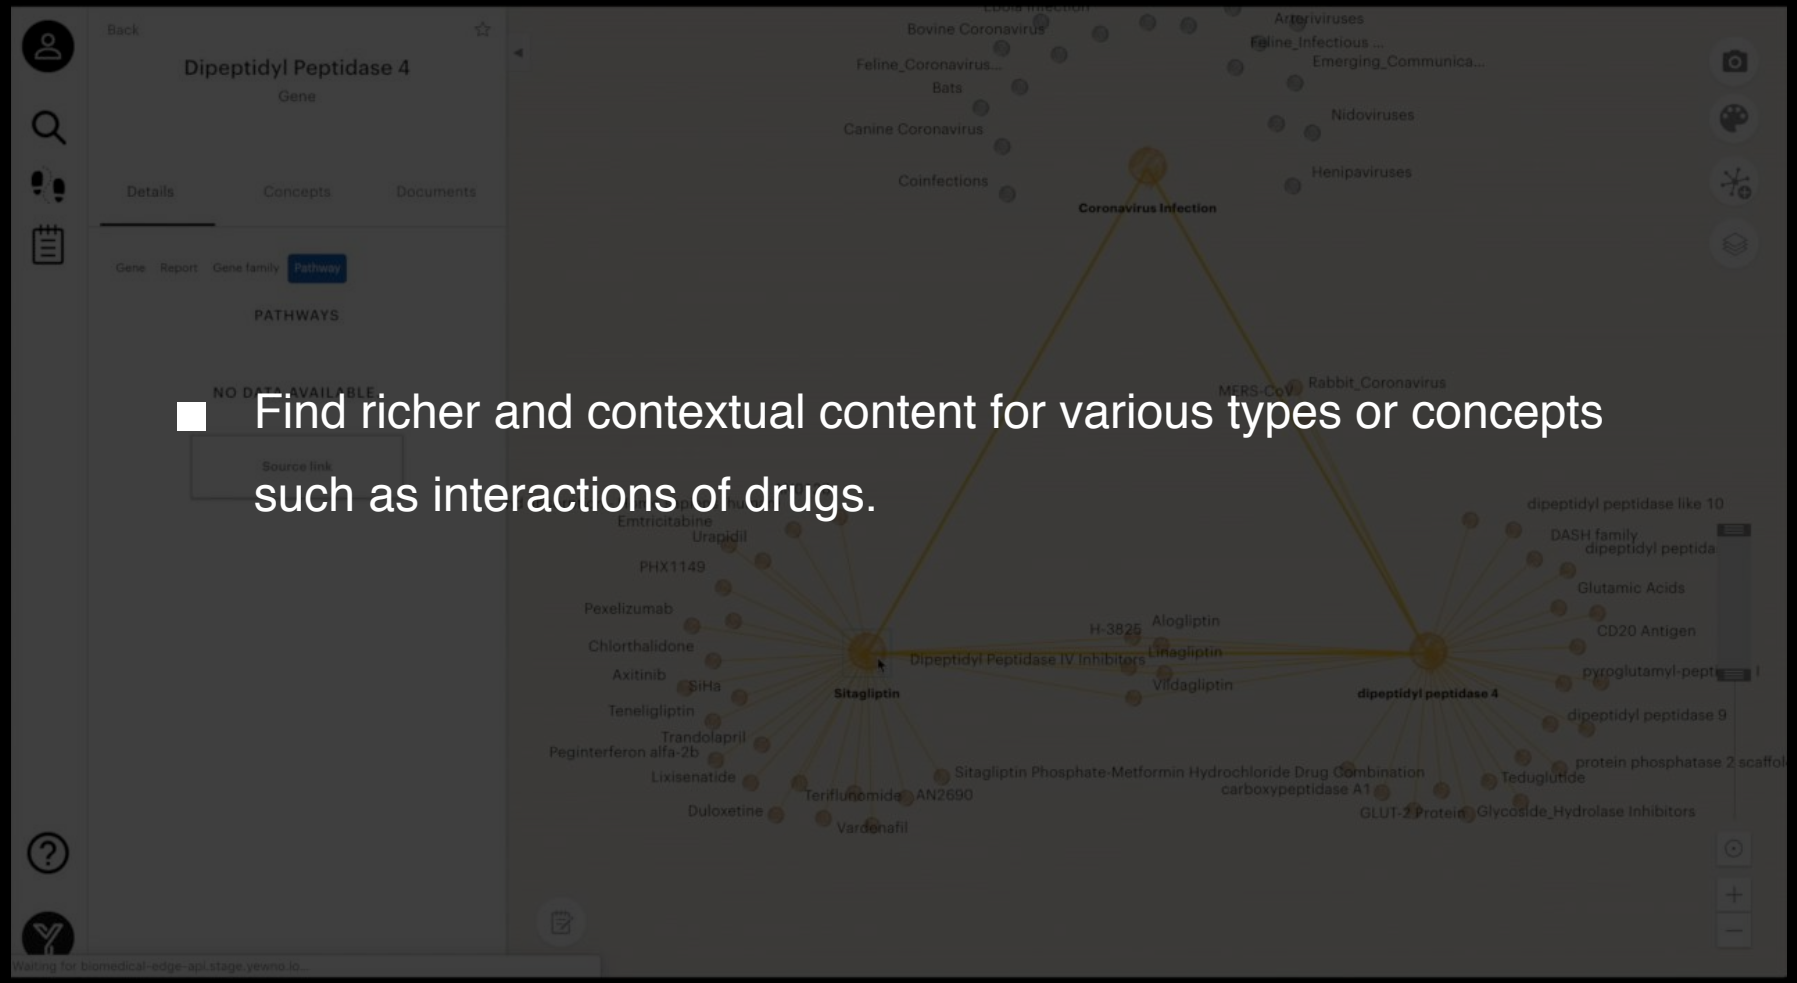

- Find richer and contextual content for various types or concepts such as interactions of drugs.

- Find richer and contextual content for various types or concepts such as interactions of drugs.

Yewno | Life Science

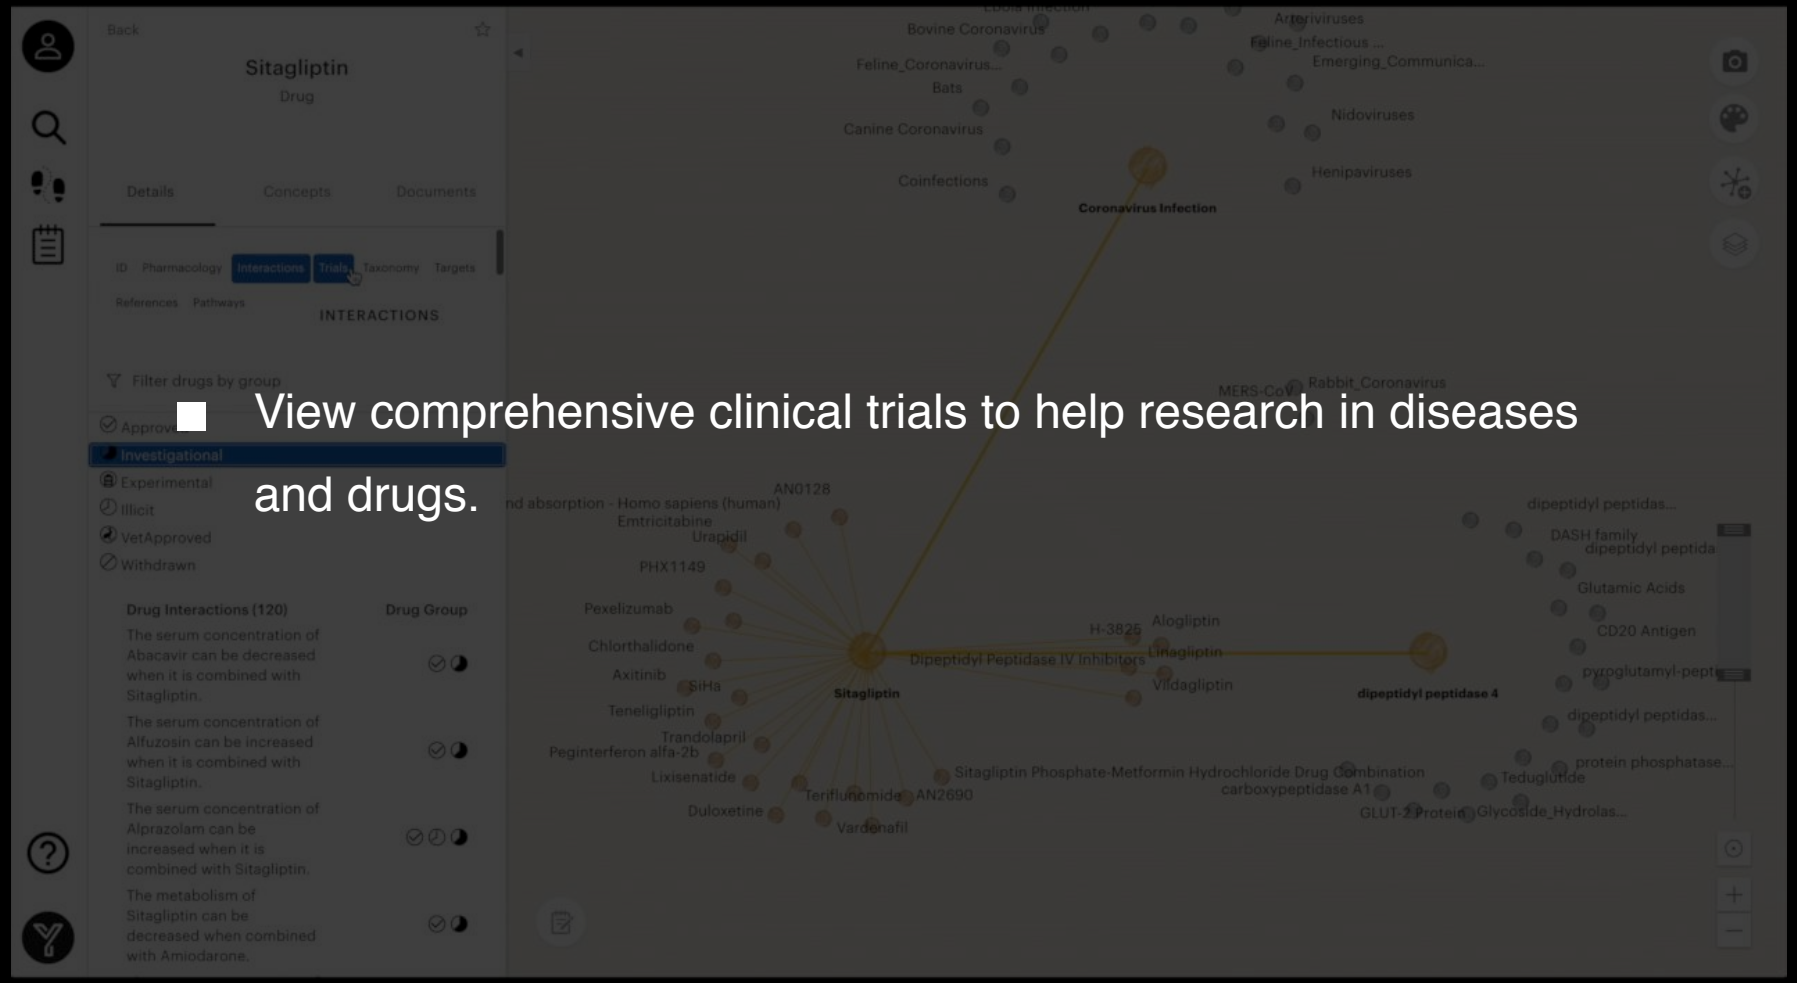

- View comprehensive clinical trials to help research in diseases and drugs.

- View comprehensive clinical trials to help research in diseases and drugs.

Finally…

How do we see the world of Knowledge now and in future?

Yewno | Open Knowledge

Preview

Machine Generated Content

THE INTRINSIC STRUCTURE OF THE LANGUAGE

In natural languages (as opposed to artificial languages), the process of forming an idea is articulated through the concatenation of processes such as:

- Sounds ⇢ Phonetics
- Phonemes ⇢ Phonology
- Words ⇢ Morphology
- Phrases ⇢ Syntax
- Meaning ⇢ Semantics
- Meaning in context ⇢ Pragmatics

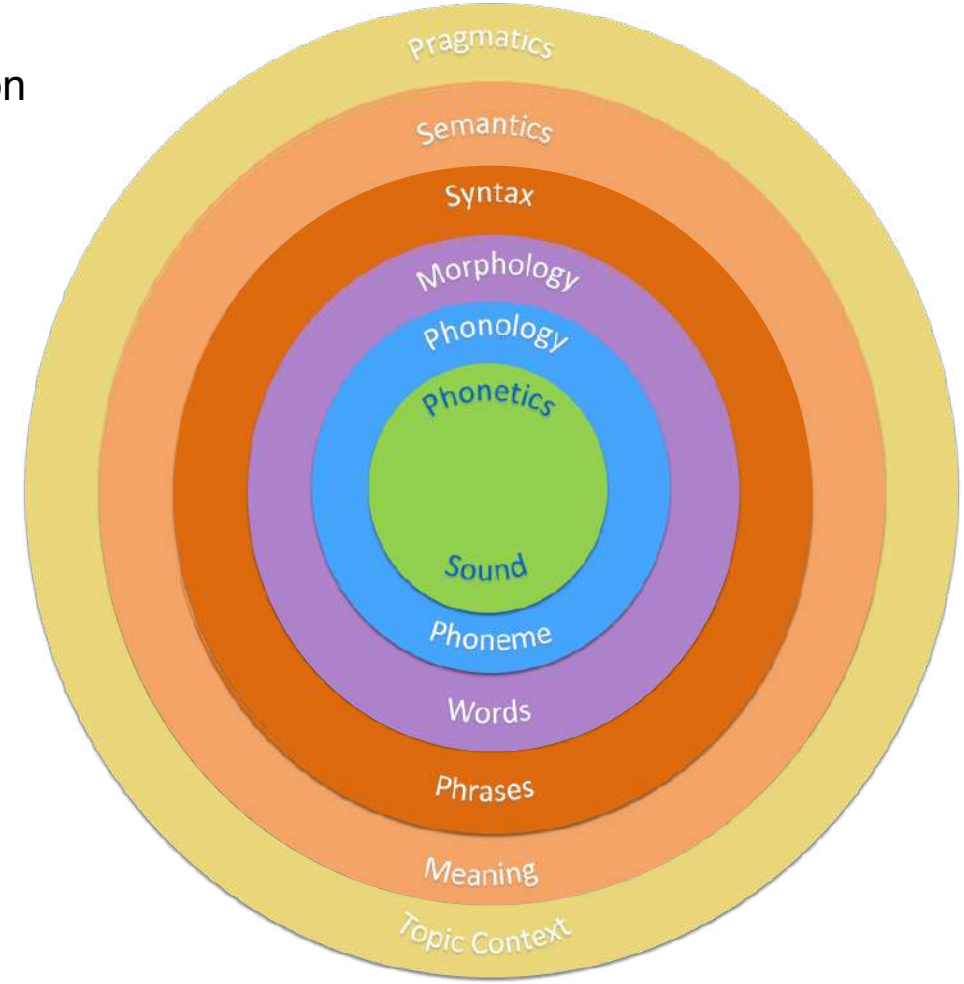

In the language:

- The syntax binds words and phrases in the representation of a concept
- Semantics, on the other hand, contextualizes the concepts within a subject

GOAL

To develop a technology for

Automatic Generation of content

Challenges…:

- What to write? (i.e. information selection)
- How to write? (i.e., render information coherently)

APPROACH

- Automatic analysis of Hierarchical Topic embedded into published material
- Identify coherent thematic areas across documents in the same topic (e.g. adversarial neural networks)
- For each thematic area generate blocks of content that summarizes through an inferential semantic structure actual research using deep learning networks

STATE OF THE ART [1/2]

Current solutions for the automatic generation of content (text) broadly fall into three categories:

1. Rule-based*: *Text is generated starting from a set of predefined production rules (i.e. if-then rules)

Pros: less sensitive to noise
Cons: output limited to the set of available rules; coverage extension requires human input

2. Planning-based*: *Text generation is seen as a classical AI planning problem

Planning is the process of identifying a sequence of one or more actions (i.e. discourse acts) to satisfy a communicative goal

Pros: allow for greater variety compared to the rule-based approach
Cons: computational costs

3. Data-based*: *Text generation is driven by a Machine Learning model trained on a (large) set of examples

Best performing algorithms today are based on either *Probabilistic Context Free Grammar* (*PCFG*) or *Deep learning*

Pros: algorithms can continuously adapt to new data
Cons: mostly black-box models and high computational demands

LIMITATIONS OF CURRENT APPROACHES

- Poor vocabulary
- Unusual structure of sentences
- Non *Zipfian* distribuiton of token frequencies
- Algorithms do not take semantic information into account

As a result, automatically generated text sounds unfamiliar to a human reader, both in style and content

How can we make it better?

Extraction of *concepts:* semantically disambiguated units of knowledge

Topic model clustering with deep hierarchical topic model

Extraction of *emerging properties* from the induced semantic graph

Identification of *semantically similar* chunks of data from big corpora

Coherent based algorithms leveraging *computational linguistics*, *machine learning* and *graph theory*

HIERARCHICAL ADVERSARIAL LANGUAGE GENERATOR

Input: set of articles, sources in a given period

Text Generation

Generative Adversarial Network for Abstractive Text Generation
(L. Liu et all, arXiv:1711.09357)

Two competing neural network models: **generator** and **discriminator**

- One takes noise as input and generates samples (the generator),

- The other model (the discriminator) receives samples from both the generator and the training data, and plays in order to be able to distinguish

The two networks play a continuous *game*

The competition will drive the generated samples to be indistinguishable from real data over time which leads to a model for text generation.

Let’s try this…..

What remains is the slender trace of lightweight watermark that marks our footsteps,

It is the idea of you that you will leave my way and I will doubtless continue my journey.

And the desolate way of winter on the bramble fields, the dim light that shines through the rows of scattered,

The gloom of time that greets the day, the mild sleep of those who are waiting for the tomorrow.

– H.A.R.I (Yewno's Hyper Associative for co-Related Inferences engine)

(Simulating Eugenio Montale)

UNDERSTANDING AND CREATIVITY OF MACHINES

Thanks to the new techniques of intelligent algorithms training, it is possible to "Teach" a machine:

- The extraction of a topic (melancholy, love, ...)
- Choosing a style
- The concatenation of concepts
- The recurring patterns grouping

Deep Recurring Neural Network (LSTM)
Coupled with Adversarial Generator/
Discriminator modules

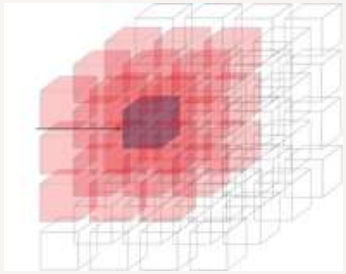

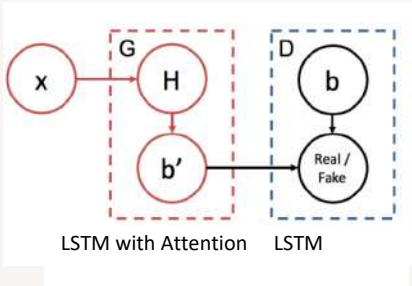

Yewno

THE FUTURE OF KNOWLEDGE IS NOW….

Yewno

Transforming Information into Knowledge