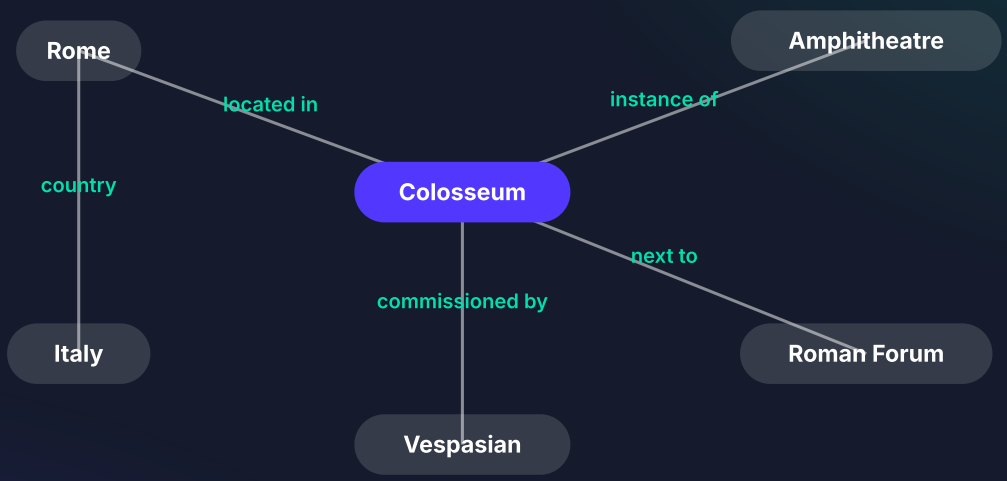
Knowledge Graphs

What a knowledge graph is, which ones shape search and AI answers, how an entity gets into them, and how to check and strengthen what they know about you.

Teaching series · Lesson 8 · intermediate

ENTITIES AS NODES, RELATIONSHIPS AS EDGES



Simplified, illustrative graph. Real knowledge graphs store each fact with an identifier and, ideally, a source.

55	5	1
minutes	graphs that matter	entity audit

Ameneh Saeednia
Senior SEO and Generative Search Strategist

Before you start

What this lesson covers

Search engines and AI systems don't just store pages. They store **facts about things**, connected to each other. That network is a knowledge graph, and it shapes knowledge panels, entity understanding and, increasingly, AI answers.

By the end of this lesson you can

- Define a **knowledge graph** in terms of nodes, edges, identifiers and sources.
- Name the graphs that matter for SEO: **Google's Knowledge Graph, Wikidata, DBpedia, your own site graph** and graphs used with LLMs.
- Read a **Wikidata item** and explain statements, qualifiers and references.
- Check what Google's Knowledge Graph knows about an entity with the **Knowledge Graph Search API**.
- Plan **realistic** steps to make your entity clearer, without promising a knowledge panel.

Before this lesson

Lesson 2 (**triples**), Lesson 5 (**entities**) and Lesson 7 (**entity schema and @id**). This lesson connects them.

How to use it

Read Parts 1 to 7, run the **entity audit on page 11**, then take the **quiz on page 12**.

Part 1

What a knowledge graph is

A knowledge graph stores facts as **entities (nodes)** connected by **relationships (edges)**. Each fact is a triple, and good graphs give every entity a **stable identifier** and every fact a **source**. [1]

Part	What it is	Example
Node (entity)	A uniquely identified thing	Colosseum, with an ID such as Wikidata Q10285
Edge (relationship)	A typed connection between two nodes	located in, commissioned by, instance of
Literal value	A plain value attached to a node	A date, a number, coordinates, a URL
Type / class	What kind of thing a node is	Colosseum is an instance of amphitheatre
Source / reference	Where the fact comes from	A cited book, an official website

Graph vs. table

A spreadsheet stores rows with fixed columns. A graph lets **any entity connect to any other**, so you can follow paths: Colosseum → in Rome → in Italy → currency euro.

Graph vs. keywords

A keyword list knows the **string** “Colosseum.” A graph knows the **thing**: what it is, where it is, who built it, and that it's different from a cinema chain with the same name.

Part 2

The graphs that matter for SEO

Several different graphs influence how search and AI systems describe an entity.

Graph	Who runs it	What to know
Google Knowledge Graph	Google	Launched in 2012. Google said in 2020 it held **500 billion facts about 5 billion entities**. Powers knowledge panels and entity understanding. [2][3]
Wikidata	Wikimedia community	**Open and editable**, data released as CC0. Every item has a Q-number; facts are statements with references. Widely reused by other systems. [4]
DBpedia	Open project	Facts **extracted from Wikipedia** infoboxes. Useful for linked-data work.
Your site graph	You	The connected schema on your pages: Organization, people, products, linked by **@id** (Lesson 7).
Graphs used with LLMs	AI developers	Language models store knowledge in their weights, not a graph, but graphs are used to **ground** answers; GraphRAG builds a graph from documents for retrieval. [5]

The connection

These graphs **borrow from each other and from the web**. A fact that's stated consistently on your site, in Wikidata, on Wikipedia and in reputable coverage is easier for every system to trust.

Part 3

Reading a Wikidata item

Wikidata is the easiest public graph to inspect. Each item has a **label, description, aliases and statements**. [4]

Field	Example: Colosseum (Q10285)	Why it matters
Label and description	“Colosseum”; a short description of what it is	The name and the **disambiguating sentence**
Aliases	Other names, including other languages	How different wordings **resolve to one entity**
instance of (P31)	amphitheatre	The entity's **type**
located in (P131)	Rome (Q220)	A **relationship** to another item
country (P17)	Italy (Q38)	Another relationship
official website (P856)	The site's URL	Connects the item to its **entity home**
Identifiers	Links to other databases	Lets systems **match records** across graphs

Statements have structure

A statement can carry **qualifiers** (when it was true, in what role) and **references** (where it comes from). Unreferenced statements are weaker.

Editing responsibly

Wikidata welcomes improvements, but follow its **notability and conflict-of-interest** guidance. Add references from independent sources; don't create items just to promote a business.

Part 4

How entities get into Google’s graph

Google doesn't publish an application form. Its Knowledge Graph draws on **many sources across the web**, and knowledge panels appear when Google is confident about an entity. [3]

1

An entity home

One page, usually **About**, that clearly states who or what the entity is, with consistent facts.

2

Consistent facts

The same name, description, founding date and leaders on your site, profiles and directories (see *Entity Clarity*).

3

Independent coverage

Reputable sources writing **about** the entity, not just mentioning it.

4

Open records

Accurate **Wikidata** and, where notable, Wikipedia entries with references.

5

Structured data

Organization or Person schema with **sameAs** to verified profiles (Lesson 7).

6

Verification

If a panel appears, the entity can **claim it** through Google and suggest corrections. [6]

What to promise

None of these steps **guarantees** a knowledge panel. They make an entity easier to identify and describe correctly. That's the honest goal, and it matters for AI answers too.

Part 5

Checking what Google knows

Google's **Knowledge Graph Search API** returns entities matching a name, with their type, a description and a relevance score. [7]

Field in the API response	What it tells you
@id (kg:/m/... or kg:/g/...)	The entity's **machine ID** in Google's graph. If your entity has one, Google has a record of it.
@type	What Google thinks it is: Organization, Person, Place, Corporation...
description	A short label such as “Travel company.”
detailedDescription	A longer description, often with a source URL such as Wikipedia.
resultScore	How strongly the result matches **your query**, relative to other results. Not a quality score.

How to run it

The API needs a **free Google Cloud API key**. Query your brand name, then variants and misspellings. Note whether the right entity appears, its ID, type and description.

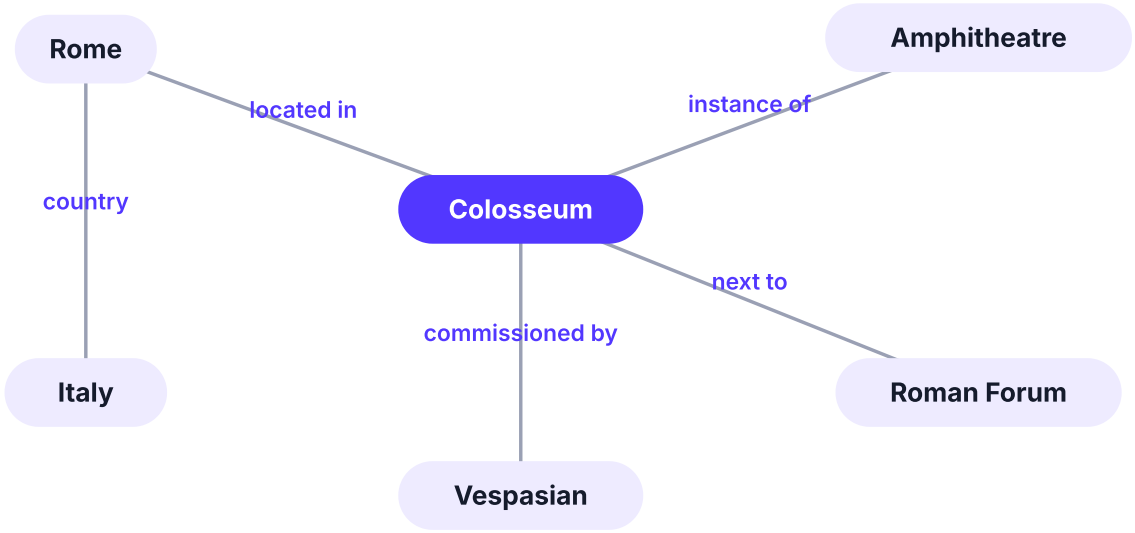
What to look for

No result: no clear entity record yet. **Wrong type**: a person read as a company. **A namesake ranks first**: you have a disambiguation problem. **Old description**: a source needs updating.

Part 6

Your site's own knowledge graph

Your schema can form a small, connected graph of your own entities. **Stable @id values** turn separate pages into one graph.



From schema to graph

Organization, people, products, places and articles become **nodes**. Properties like **author, provider, about, location** become edges. Define each node once and reference it everywhere.

Internal links are edges too

A guide page linking to the Colosseum tour, and the tour linking back to the guide, mirrors the graph for **readers and crawlers**. Keep links and schema telling the same story.

Part 7

Knowledge graphs and AI answers

LLMs generate text from patterns learned in training. Graphs help in three ways, and they also show why consistency matters.

1

Grounding

Systems can check or supplement an answer with **structured facts** from a graph, which reduces made-up details.

2

Disambiguation

A graph tells a system **which Paris** or which Apple is meant, using connected entities.

3

Retrieval

GraphRAG and similar methods build a **graph from documents** to answer broad questions across many sources. [5]

Why your facts need to agree

If your site, your profiles and Wikidata give **different founding dates or descriptions**, a system has to choose, and it may choose wrong. Consistent, referenced facts reduce that risk.

What we don't know

AI products don't publish exactly which graphs they use or how. Treat graph work as **improving the evidence** about your entity, not as a direct AI ranking lever.

Common mistakes

What doesn’t work

1

Chasing the panel

Treating a knowledge panel as the goal. The goal is **accurate identification**; panels may follow.

2

Promotional Wikidata items

Creating items without independent references. They can be **deleted** and damage trust.

3

Inconsistent facts

Different names, dates or descriptions across sources. **Pick one canonical version** and update everywhere.

4

sameAs to anything

Linking to pages that aren't the **same entity**. That merges records incorrectly.

5

Ignoring namesakes

Not checking who else shares your name. **Disambiguating context** has to be on your pages.

6

Expecting fast change

Graphs update on their own schedule. **Recheck quarterly**, not daily.

Exercise

Audit one entity across the graphs

Time: 25 minutes. Use your company, your personal brand or a client.

Step	What to do
1	Write a **canonical fact sheet**: name, description, type, founding date, location, founders or leaders, official site.
2	Query the **Knowledge Graph Search API** for the name. Record the @id, type and description, if any.
3	Search **Wikidata** for the entity. Record the Q-number and check each statement against your fact sheet.
4	Search Google for the name. Note whether a **knowledge panel** appears and what it says.
5	Check your **About page and Organization schema**: same facts, sameAs to verified profiles, one @id.
6	List every **mismatch**, the source that holds it and who can fix it. Recheck in three months.

What good looks like

Every source states **the same core facts**. The right entity appears for your name, with the right type. Mismatches have **owners and dates**.

Discussion question

Which single source, if corrected, would fix **the most** downstream errors? Usually it's the entity home or Wikidata.

Check your understanding

Quick quiz

1. In a knowledge graph, a relationship like “located in” is:

a) A node
b) An edge
c) A literal value

2. Wikidata items are identified by:

a) Q-numbers
b) Keywords
c) Page URLs

3. In the Knowledge Graph Search API, resultScore means:

a) How trustworthy the entity is
b) How well the result matches your query
c) The entity's ranking in Search

4. Which step can guarantee a knowledge panel?

a) Organization schema
b) A Wikidata item
c) None of them

5. Why do consistent facts across sources matter for AI answers?

a) Systems may pick the wrong version when sources disagree
b) Consistency is a ranking factor
c) AI products publish their sources

6. Stable @id values in your schema:

a) Turn separate pages into one connected graph
b) Create a knowledge panel
c) Replace sameAs

Answers and sources

Answer key, glossary and further reading

Answer key

1: **b**.
2: **a**.
3: **b**.
4: **c**.
5: **a**.
6: **a**.

Glossary

Node: an entity in a graph.
Edge: a typed relationship between nodes.
Q-number: a Wikidata item ID.
Entity home: the page that best defines an entity.
Grounding: checking answers against sources.
GraphRAG: retrieval using a graph built from documents.

01 **Knowledge Graphs. Hogan et al., ACM Computing Surveys, 2021**
A thorough academic introduction.
https://arxiv.org/abs/2003.02320

02 **Introducing the Knowledge Graph: things, not strings. Google, 2012**
https://blog.google/products/search/introducing-knowledge-graph-things-not/

03 **A reintroduction to our Knowledge Graph and knowledge panels. Google, 2020**
Scale of the graph and where knowledge panel information comes from.
https://blog.google/products/search/about-knowledge-graph-and-knowledge-panels/

04 **Wikidata: Introduction. Wikidata**
https://www.wikidata.org/wiki/Wikidata:Introduction

05 **From Local to Global: A Graph RAG Approach to Query-Focused Summarization. Edge et al., 2024**
https://arxiv.org/abs/2404.16130

06 **Get verified on Google and claim a knowledge panel. Google Help**
https://support.google.com/knowledgepanel/answer/7534902

07 **Knowledge Graph Search API. Google for Developers**
https://developers.google.com/knowledge-graph

Lesson by Ameneh Saeednia, prepared with AI assistance. The Colosseum graph is simplified for teaching. Related reading on seogirl.ca: *How Google Actually Mines Wikipedia for Knowledge Graph Facts* and *The Knowledge Graph Didn't Retire When AI Search Showed Up*.