From a General Interest to the State of the Art using AI Tools – Guiding Notes

Leonel Morgado

V. 1.1, Apr. 9, 2025, translated by ChatGPT on Apr. 10, and reviewed by the author on Apr. 29.

Abstract
This document is intended to support students in the initial stages of research, particularly in courses dedicated to the planning of master's dissertations and doctoral theses. It proposes that identifying the state of the art in a problem (and its associated techniques) be understood as a dual process: technical and personal. On the technical level, it involves understanding the current point of knowledge (what advances have been made, which techniques are in use, and how to apply them); on the personal level, the focus is on exploring alternatives and leading the student-researcher to reflect on their priorities, ambitions, and constraints. The goal is to empower them to make informed and conscious choices about the contours of the problem and the technical-scientific goals to pursue.

1. Introduction

Hey there, student! You've come here seeking guidance on how to conduct a literature review or a survey of the state of the art. But with what mindset? Are you feeling amidst an exploration, excited about what you'll discover? Or are you facing a wall, struggling to decide whether to jump over or tear it down to move forward? Whatever your mood, this is the beginning of a journey.

Those feelings can be contradictory. Enthusiasm gives us energy to move ahead, but also brings insecurity, when faced with the vastness of knowledge—so immense, it's impossible to fully digest, as if it's going to drown us... expose us as perpetual ignoramuses, frauds merely splashing on the surface of an ocean. How could we possibly become masters or doctors with such ignorance? The struggle gives us energy and determination to finish, but it can also blind us, keeping us from seeing what we're really doing: what we lose or disregard in our rush to just get the obstacle out of the way. What mental confusion do we suffer as a result, when we feel overwhelmed by so many contradictions, alternatives, and details?

Comfort will come from the awareness that **conscious ignorance is a sign of deeper knowledge**. To know more is to accept how little we actually know. Permanent ignorance is the humble state of all researchers: they simply set out on the path and took enough steps to see that the world was much bigger than it seemed. This document aims to help you gain a notion of the "state of the art" in your topic. To help you become aware of that vastness, so you have the freedom to decide. To reach that awareness, you will go through a review of the technical-scientific literature.

As in an adventure movie, the main character ventures through paths where previous explorers sometimes met their doom… We want to make sure you're not one of those "victims." The victims of the literature review are those who thought it was just meant to make the final document look good. Those that prepare it like a courtroom defence against imaginary accusers of laziness or ignorance: "Your honours, I read! See? I know! Here's the proof!" In that kind of defence effort, which aims to hide and protect… one ends up almost empty. With some ideas, of course: the effort was made, time was spent… but little was gained. You end up not much better than when you started, just more tired and drained.

Conducting a literature review is not merely a technical process. It is the combination of two interconnected processes: the technical and the personal, which feed into each other.

At the technical level, the literature review seems very objective: a rigorous search for knowledge. In a somewhat naïve or idealized view, it all looks methodical: you define a goal, identify the right keywords, search for relevant articles and authors, analyse them, extract the key points. And at the end, like a long-lost treasure… you're left with the current knowledge on your problem… if only it were that simple and objective!

But the literature review, the quest for the state of the art, is also an inner journey. As you encounter and interpret relevant works, you'll see that there are many ways to analyse, explore, and develop the topic. More than you can possibly pursue. And then dilemmas arise that stall the technical process:

- What makes sense to you, what are your priorities, interests, and ambitions?
- What do your time, resources, and team constraints suggest?

This is the human and transformative side of the review: made out of doubts, choices, and gradual self-awareness. The review is no longer just about what's been done or said; it becomes a pilgrimage to discover what we want to do and who we want to be in our research.

It is the balance, the ongoing dance between technical actions and personal decisions that will give meaning to this path. Neglecting the methodical rigor of the technical and procedural aspects leads to empty chatter, pages filled with hollow text… which may swell us with false confidence, but the ignorance never goes away, it just hides, tail peeking out. Ignoring the personal aspects will result in sacrificing our future efforts to a mechanical result that doesn't reflect who we are. **This dance between the technical and the personal builds something more beautiful: a clear purpose and a roadmap to reach it**. And with those two elements, you can build research that is both solid in method and meaningful in contribution.

2. The General Process

This dance between the technical aspects and the personal aspects will unfold throughout a general guiding process (Figure 1). Each step assumes the need to find options, alternatives, and clarity. It recognizes the importance of bringing our concerns and intentions to the surface, to clarify the next steps.

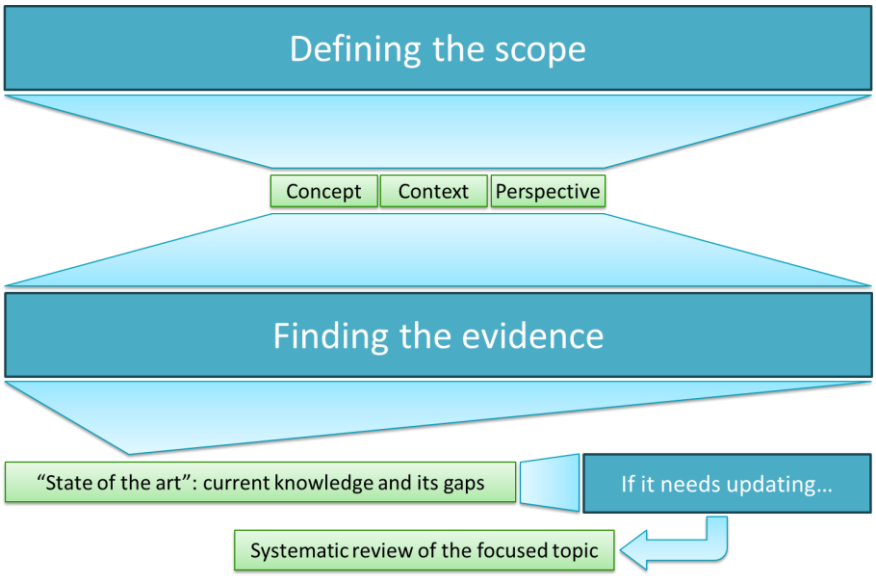

Figure 1 – The general process

The following sections of this document will present these steps in detail, but this is the general overview. We start with an initial idea, which may already *seem* focused, or it may be just a general area of interest.

The first step aims to define the scope of the review, our focused topic: to specify a more concrete objective, reflected in three parts: a concept, its context, and the perspective through which we explore it.

The second step aims to find evidence to support that objective. In a systematic way, we refer to what other explorers—those who have returned from the unknown—have already synthesized, to get a preliminary sense of what we aspire to understand: the "state of the art", that is, what humanity already knows and does not know about our objective.

- **What is already known** keeps us from reinventing the wheel, tells us how we can proceed effectively and efficiently, and warns us of pitfalls where others have already stumbled.
- **What is not yet known** helps us imagine our contribution, chart paths and plans… and grasp the potential value and impact of our effort.

This division isn't as clear-cut as it may seem in this simplified presentation: the "known" is tied to assumptions, constraints, and degrees of confidence. The "unknown" is not always explicitly expressed: it's the "known unknown," when someone has already considered doubts, data gaps, weak assumptions, and so on. But most of what is unknown… has yet to be imagined: it is the unknown we are unaware of. For example, it is the uncertainty or doubt that arises when we try to look at reality through a new inspiration. It is the effort to imagine what might not be visible.

This reflection—both technical and personal—on the state of the art may then lead us to the decision to proceed (or not) to **the third step**: a systematic review of our focused topic. It requires a personal judgment on the risks and benefits of proceeding solely based on the state of the art gathered in the second step. If we continue with techniques and processes that have already proven inadequate, we may end up wasting our effort or even invalidating our results. If we explore a topic only to find it had already been explored in the same way, our contribution is limited to verifying the reproducibility of existing knowledge.

The most dangerous risk, however, is missing the true value of the gem we may be about to uncover with our effort:

— Having a diamond but mistaking it for an ordinary stone. This happens when we're so attached to an initial idea, and so unaware of the broader picture, that we are unable to see our work in a different light. The history of science is full of discoveries born out of chance (Copeland, 2019). That mould that spoils our samples might actually be penicillin, the first antibiotic! Those 100 failed algorithms that don't optimize the problem… knowing that is valuable—it could save other researchers a lot of time.

To avoid discarding diamonds, we must be prepared to recognize them. Being well aware of what is known and unknown in a field helps us recognize the value of what arises, even if it seems irrelevant or even seems like an obstacle. Being prepared requires more than goodwill and superficial knowledge: we must be able to draw connections and recognize patterns. A solid understanding of the state of the art gives us more opportunities to make these connections and detect those patterns, so we can benefit from serendipity or apparent setbacks.

— Having a diamond but being unable to explain it to anyone. This is also common, especially in creative fields like engineering and the arts. Our discoveries and creations are personal, intimate. Their value is obvious to us. But… what about to others? If our discovery or creation is not a direct improvement or extension of what exists, its value may be hard to communicate. Without an understanding of the many different ways other people view the world, and what they find simple or

obvious… it will be difficult to help them see our work the way we do. Especially if people have to go through unfamiliar processes, contexts, or situations: and even more so if we're asking them to imagine something that doesn't yet exist.

Here again, a strong sense of the state of the art plays a crucial role: it provides us with the conceptual and communicative tools to present our discoveries or creations in a comprehensible context. The state of the art helps us structure the narrative, building bridges between our work and others' understanding. Without that communicative foundation, even the brightest diamonds can seem like rough stones.

Recognizing and valuing what we do—and being able to communicate it effectively—requires something beyond intuition or brilliance. It requires preparation, reflection, and awareness of the existing landscape and future possibilities. That way, we avoid wasting brilliant ideas, turning them instead into contributions that not only shine for us, but also illuminate the way for others.

Finally, **that communication is what allows science… to be science**. If we can only present our ideas without debate, seeking followers, then it's not science: it's doctrine. We need to subject our ideas to external scrutiny, allow our arguments to be challenged, and let our steps be replicated, so they can be confirmed or disproven. And if we fail to communicate effectively, those processes may never begin… or be rendered impossible.

3. First Step: Defining the Scope

3.1. The Iterative Cycle of Exploration

To begin this step, we must recognize—and accept—that our topic is too vague, too unspecific. It might not seem that way, so dismantling that illusion is a personal task we must undertake. This step will consist of discovering the research terrain and the possibilities it offers.

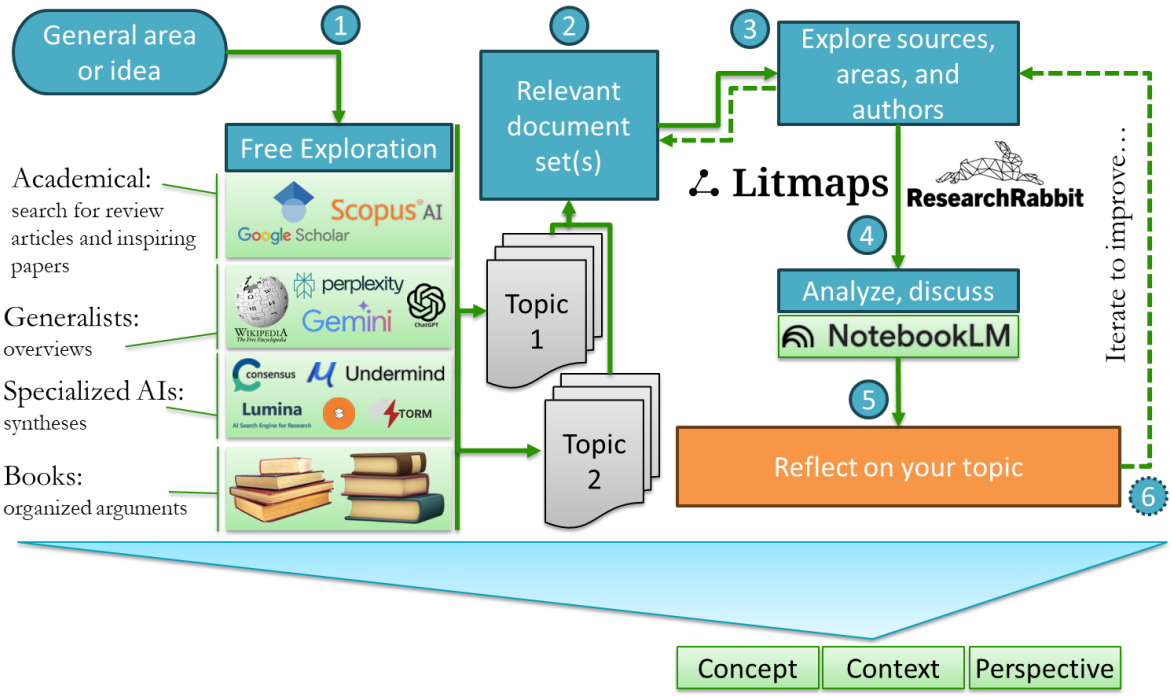

Figure 2 – The scope definition step

It consists of six steps, in a cycle (Figure 2):

1. **Free exploration** of the area or general interest.
2. **Organizing** one or more sets of relevant documents.
3. **Exploring** sources, areas, and authors from those document sets, enriching them.
4. **Analysing and discussing** those document sets.
5. **Reflecting** to synthesize the topic into a focused concept, context, and perspective.
6. **Iterating** through the previous steps to refine or redirect that focus.

3.2. Step 1: Free Exploration

Free exploration of the area (or general interest) aims to clear the fog from the landscape, to give us a sense of possible paths or destinations. Each word, each concept, may hide different or even contradictory interpretations and meanings: sometimes subtly, sometimes in a bewildering way. At this stage, our lack of familiarity with the terrain is at its peak: we're just starting out, and trying to systematize this vastness would be foolish! If we embrace randomness, the richness of wandering, the persistence of constantly reflecting to reconsider our goals… we'll be continuously clearing the fog.

Figure 2 suggests four exploration routes, but the emphasis is on the word "free": others may come to you! These options and the logic behind each are shown to encourage reflection on their distinct contributions, helping us see new perspectives on the unknown. Repeating similar processes tends to give us the same perspective, so it's valuable to notice when our processes are genuinely different.

3.2.1. First Route: Searching for Academic Review Articles

In most areas or topics, if we just run a general search, we'll get a flood of results. This issue is not new and has long been addressed by producing overview works of the existing literature—so-called *panoramic studies*. These fall into two major types (Kitchenham et al., 2016, p. 32):

- **Literature Reviews**, which summarize the search and reading of published works on a topic. Since they are works about other works, they are considered "secondary studies." They seek to summarize results, problems, methods, techniques, etc. When they attempt to reconcile divergent findings and detect patterns, they are called *meta-analyses*. If many reviews already exist in a field, we might find works that compare and summarize them: *reviews of reviews*, known as "tertiary studies." Common terms in the titles of such reviews are *review*, *survey*, and *meta-analysis*.
- **Literature Mappings**, which aim to identify *where* information exists on a topic. A review seeks to summarize; a mapping focuses on identifying what's out there: something common in areas where it's hard to compare reports or cases, or where understanding information gaps is important. These are also secondary studies. Like reviews, mappings can also be conducted over other panoramic works, making them tertiary studies. Common title terms include *review of review*, *survey of surveys*, and *mapping study*.

So, if during your free exploration you don't just search for scattered papers, but instead look specifically for panoramic studies, you can already benefit from a structured understanding of the field: its problems, results, approaches, techniques, interpretative theories, and other potentially relevant aspects.

Although academic search engines sometimes allow you to specify that you're looking for review articles, how well this works varies a lot. For now, the most effective way is to explicitly indicate you

want to find works with *"review"*, *"survey"*, *"meta-analysis"*, or *"mapping study"* in their titles—these are the most common title keywords for such works.

Each search tool has its own way of doing this. One of the most widely used is Google Scholar (scholar.google.com), which tends to have exceptionally broad coverage of academic literature (Martín-Martín et al., 2021), though not complete (Gerasimov et al., 2024). You can indicate you want these terms in the title using quotes for exact match, "intitle:" to specify the field, and "OR" or "|" for alternatives, for example:

```
(intitle:"review" | intitle:"survey")
```

Keep in mind that having these words in the title doesn't guarantee the result is a panoramic literature work. Some articles with these words are in fact book reviews ("review"), field surveys ("survey"), or other types of work.

Also, besides Google Scholar's incomplete coverage, another issue is the signal-to-noise ratio. That is: even among panoramic works, many may be low quality, wasting time and requiring experience to evaluate. One strategy to address this is to limit the search to well-regarded publication venues. This doesn't guarantee quality (good work may appear in unknown venues, and bad work in reputable ones), but it tends to improve average quality. A common choice for this is **Scopus** (www.scopus.com), which limits searches to publications indexed in that scientific database.

3.2.1. Second Route: Consulting Generalist Secondary Sources

When beginning to explore a topic, it's natural to feel prematurely narrowed in scope. For example, we might only know one or two ways to talk about the subject, just two or three areas of interest or analysis… A quick way to broaden our horizons and break this initial freeze is to consult generalist sources, such as encyclopaedias and dictionaries (e.g., Wikipedia) or interact with general-purpose artificial intelligences (AI), such as OpenAI's ChatGPT, Mistral's LeChat, Google's Gemini, Microsoft's Copilot, or Perplexity, among many others.

These generalist sources should not be considered foundational, but rather as tools for ice-breaking or unlocking ideas. For instance, they can help us see what different themes are associated with the area, what terminology is used, and what aspects are mentioned. They may lead us to new lines of thought we hadn't previously considered, or allow us to debate them.

The fundamental limitations of secondary sources (especially generalist ones) are that they tend to rehash commonplaces about the subject, frequently falling into misconceptions or speculation that more specialized sources avoid or explicitly warn about. In the case of AI tools, they are also prone to hallucinations or confabulations.

3.2.2. Third Route: Consulting Specialized AI Services

The review of technical-scientific literature, given its nature, is a particularly suitable case for summaries produced by techniques that analyse unstructured information. Naturally, there has been a flurry of weekly launches of specialized tools for this purpose. By way of example, here are a few:

Scopus AI: Integrated into the main Scopus platform. It allows interaction with indexed scientific works, such as identifying emerging topics, providing summaries, suggesting research questions, generating lists of key documents, concept maps, and more.

Consensus: Designed to help formulate research questions and get answers based on scientific articles. It analyses available evidence, highlighting trends and viewpoints in the literature to help synthesize information quickly and reliably.

SciSpace: Performs searches and article summaries, analysing and trying to synthesize results, with a set of additional mini-tools like "PDF to video."

Lumina: Searches scientific literature with the ability to filter by year, journal ranking (quartile), etc., through queries rather than traditional search fields. It synthesizes results in text and suggests follow-up questions.

Undermind: Explores scientific articles to find concept associations, identify gaps in the literature, and suggest new research angles based on citation patterns and emerging topics.

STORM e Co-STORM: Pre-writing synthesis tools that combine searches from multiple perspectives (STORM) or simulated dialogue with experts (Co-STORM).

In a rapidly evolving field, this list is merely indicative, and the descriptions may already be outdated by the time you read this. Take the initiative to stay informed and try new tools, given how complex and important literature synthesis activities are.

A fundamental difference between these services and general-purpose AIs lies in the **reliability of sources**: they are based exclusively on **primary sources**, meaning so-called "academic" publications (the term “academic” is traditional, even though they’re not always produced within academia).

The specific features of each of these tools vary widely. Most provide area summaries, but the ways they support interpretation or further exploration differ. For example:

- Some do more than just list the sources used in their summaries; they also indicate degrees of relevance or closeness to your topic.
- Others suggest research questions or follow-up directions based on the findings.
- Still others present the key elements found in each article they considered relevant.

Some services go beyond merely listing the sources they used for the syntheses: they also indicate **degrees of affinity** between those sources and the topic in question. Others suggest **follow-up questions** or directions based on the results. Still others provide the **key aspects** they identified in each article deemed relevant.

It is important to be aware that **there are often subtle differences in terminology** between sources, which can only be inferred through a more careful analysis of their meaning and scope. An automatic synthesis, just like a human one, can fall into such misunderstandings, it is no guarantee of accuracy.

The reason we cannot consider that these services provide us with a **ready-to-use state of the art** for our problem stems from the **dual nature** mentioned in the introduction: the discovery of the state of the art is not just a technical task: it is a **dance**, in which we rediscover, step by step, our priorities, our focuses and constraints, our purposes. Intelligent syntheses had to use certain criteria to select, organize, and prioritize the disconnected information. Those criteria will not necessarily be our own, especially at this exploratory stage, when we **don’t yet know what our own criteria are**!

Often, it is precisely during the exploration process that we discover **we had been assuming or intuitively adopting aspects that are not irrelevant** and require reflection and decision-making. Therefore, the syntheses should be **part of this process of discovery and reflection**, so that we can learn how to guide them, refine them, and reorient them.

3.2.3. Fourth Route: Books

Finally, a step into the past, which is also the present that never left. Books, whether printed or digital, offer spaces for **consolidation** in an environment that is generally chaotic and disjointed. Just as

generalist secondary sources allow us to broaden our horizons, a book presents us with **structured information** about a field, more deeply than those generalist sources do.

That structure is naturally **biased** by the author's vision. Unlike generalist sources, which are often anonymous or semi-anonymous, it is usually possible to find short bios of book authors, either within the books themselves, on publishers' websites, or through professional platforms like LinkedIn. We can therefore use them as a more in-depth counterpoint to generalist sources, with attributable authorship.

The degree of reliability is very variable, and the timeliness is usually lower than in scientific articles. However, we have access to online spaces dedicated to reviewing and critiquing books, such as academic book reviews and other evaluative sources, which help us assess the relevance, credibility, and trustworthiness of these works.

Thus, we can complement the more synthetic views from the other routes with more comprehensive perspectives that help structure our understanding of a topic.

3.3. Step 2: Organizing Relevant Document Sets

The result of the exploratory routes mentioned above will be a collection of notes, annotations, and ideas, but also of **primary sources** that support them: books, journal articles, conference proceedings, repositories, etc. There will also be an indirect outcome: the realization that **complementary areas** are relevant, even if they don't necessarily intersect.

For example, imagine you've been exploring the topic of **immersive learning**. You might find that, in addition to works focused specifically on that concept ("immersive learning"), there are also studies about the **phenomenon of immersion** that are not necessarily tied to learning contexts; and studies about the **learning process** you wish to explore, but not necessarily from the perspective of immersion. In this case, you might not end up with just one set of relevant documents, but two or even three sets, each dedicated to these different aspects of the topic. Some works may appear in multiple sets: they don't have to be mutually exclusive.

These sets should be **organized for later processing**. More than classic solutions like file folders for each theme or topic lists, you should take advantage of modern support systems, such as:

- **Reference managers** (e.g., Zotero)
- **Knowledge managers**, which allow you to register and interlink notes on those documents (e.g., Obsidian)

Favor open solutions, which let you explore these sets of documents and notes in other tools, such as those used in the next step. For instance:

- You can export Zotero collections in various bibliographic reference formats, the most common being BibTeX. https://www.bibtex.org/
- Obsidian saves notes as plain text files in markdown syntax, allowing you to work with them in virtually all text analysis and processing tools—including AI tools. https://en.wikipedia.org/wiki/Markdown

3.4. Step 3: Exploring Sources, Areas, and Authors of These Document Sets

From the document sets you've compiled, you can use several techniques to identify relationships among them, helping to illuminate the informational structure of the field. There are three main aspects to consider:

1. **Which works cite or reference each other?** When one work refers to another, it has been influenced by it (even if only to disagree). There's some noise in this relationship (authors sometimes cite works they haven't read, due to social pressure, carelessness, fraud, etc.) but in general, there is more signal than noise in citation networks. A highly cited work is usually influential, so its meanings, terminology, and conclusions tend to echo throughout other works in the area. A work that cites many others tends to be aware of the topics in those references, while one that cites few or none often lacks that awareness. By reading them, you can detect biases or perspectives they ignore.
2. **Which authors work together?** By identifying patterns of co-authorship or regular citation between authors, we can detect conceptual funnels or habits. For example, imagine a group of authors uses "salmon" and "orange" as synonyms, while another group treats them as distinct colours. Recognizing clusters of co-authorship makes it easier to spot differences in language or theoretical frameworks. We may also notice that although most works focus on certain topics, less explored topics are developed by smaller author groups.
3. **What thematic areas of concern exist within a topic?** By looking at emerging citation and co-authorship networks, as well as other relationships such as themes addressed, keywords, etc., we can detect clusters of subtopics, approaches, and concerns that might otherwise go unnoticed.

As brief examples, two tools worth exploring are ResearchRabbit and Litmaps. You may also explore others like ConnectedPapers or CiteSpace. Many more will likely have emerged since these lines were written. The results from this exploration will allow you to enrich, reformulate, or reorganize your document sets:

- **You may identify foundational, influential articles**, cited by several of the documents you had already selected.
- **You may find more recent articles** that haven't yet influenced the field but intersect your interests in meaningful ways.
- **You may decide to split a set into subsets** because you've realized that distinct perspectives deserve to be considered separately.

Starting with ResearchRabbit, Figure 3 shows an example result after feeding this tool with a set of 53 documents (visible on the far left of the image). In this case, the seeded articles were about learning with virtual, augmented, or mixed reality systems, dated between 2020 and 2022. In the second column, ResearchRabbit shows that it found 1,614 documents it considers similar ("Similar work"). Clicking on that option brings up a list of those documents (which you can filter or sort in various ways) and a graph-style visual representation of the connections among them. Note that the graph only shows some of the 1,614 results, not all—so filtering and sorting can immediately help you gain new insights.

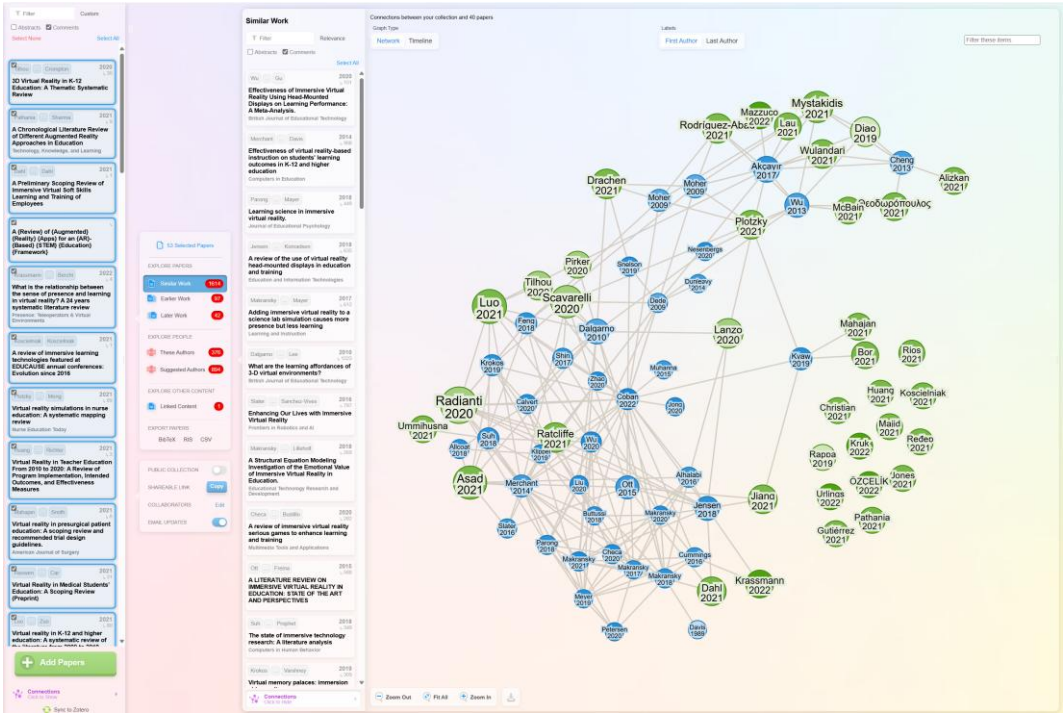

Figure 3 – ResearchRabbit results for "Similar work": articles related to the provided set

In the graph, the originally provided articles are shown in green, and the similar ones in blue. Based on the visual layout that emerges, we can immediately detect distinct clusters, which I marked with ellipses in Figure 4. This allows us to probe these different clusters directly, instead of scrolling endlessly through a long list trying to find meaningful distinctions.

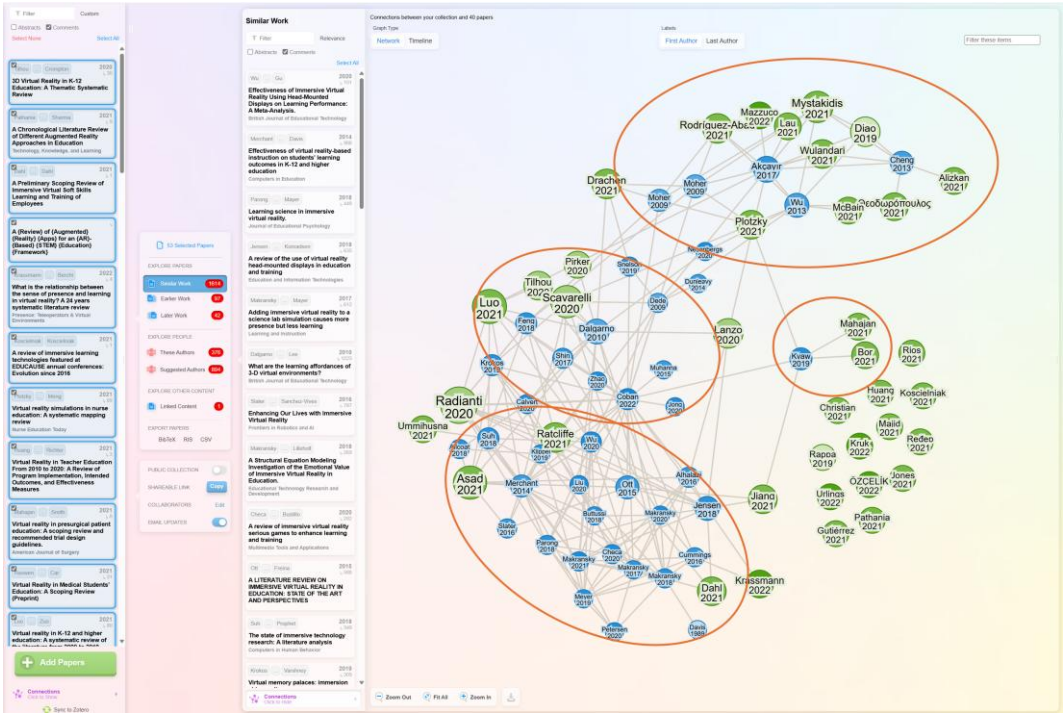

Figure 4 – Apparent clusters of suggested articles (in blue)

By clicking on the nodes in this graph, a side panel opens with details about each article, giving us a sense of its theme. Starting with the smallest cluster, containing a single article, we see that it focuses on virtual reality in health professional education (Figure 5).

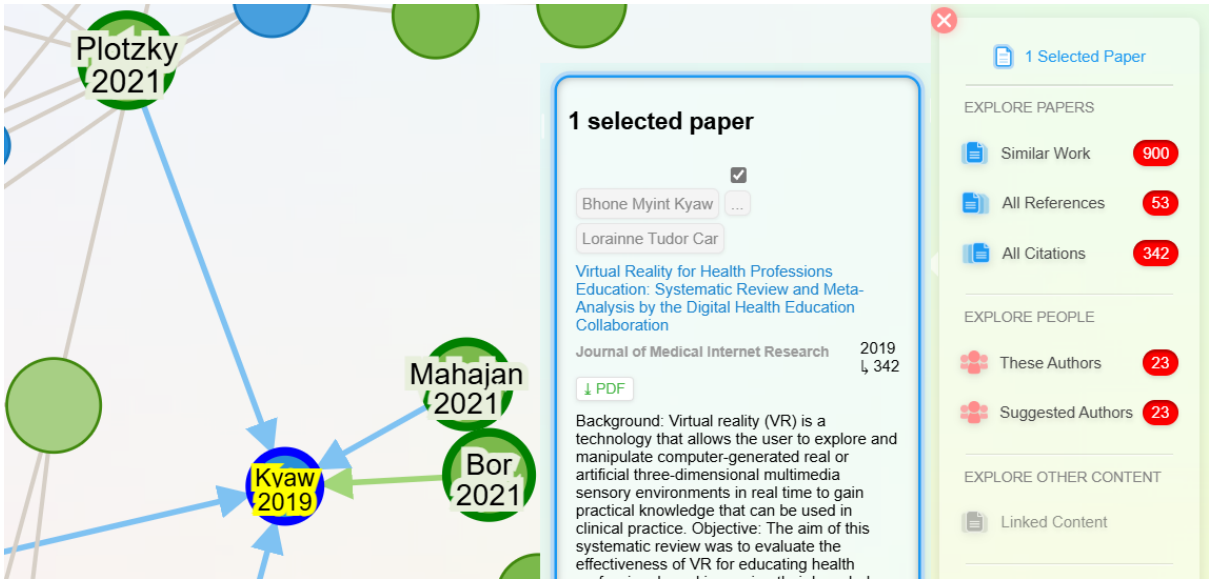

Figure 5 – Article from one of the detected clusters

Looking at the other clusters' central articles, we get the following picture:

- **The right-hand cluster** deals with **virtual reality in healthcare education**.
- **The top cluster**, moving toward its centre, features articles on **augmented reality in general education**.
- **The central-left cluster** focuses on **functionalities for learning in virtual reality**.
- **The bottom cluster**, again at its centre, gathers articles on the **effects of virtual reality on learning**.

So, right away, we identify **four thematic areas** that prompt us to think about our intent—how deeply we want to explore each of them: augmented reality (top), virtual reality in terms of functionality (left) and in terms of impact (bottom), and health-specific concerns (right).

Note that we're only referring to the *central articles* of the apparent clusters. If we move further out from the centre, we find additional focus areas. For instance, the article furthest to the top-right also addresses **augmented reality**, but now **specifically in science education**, not general education.

The key takeaway is that this **spatial organization** helps us become aware of the existence of subdomains and distinct perspectives, which can **guide our next steps**.

As you can see in Figure 5, ResearchRabbit also tells us there are:

- **900 related articles**,
- **53 references** cited by the current article,
- **342 citations** it received,
- **23 other works by the same authors**, and
- **23 suggested authors**, based on relationships and shared characteristics.

Any of these pathways—citations, co-authors, similar works—can pull us deeper "down the rabbit hole". Clicking on these opens the corresponding lists, generating new graphs of relationships. As shown in **Figure 6**, ResearchRabbit provides links to the **900 works** considered "similar" to the one that caught our attention. Apparent clusters in that new graph will lead us to:

- More works aligned with the same line of virtual reality in health professional education,
- Perspectives such as criteria for systematic review methodologies,

- E-learning for health professionals as a broader conceptual category,
- Or specific application areas, such as anatomy education.

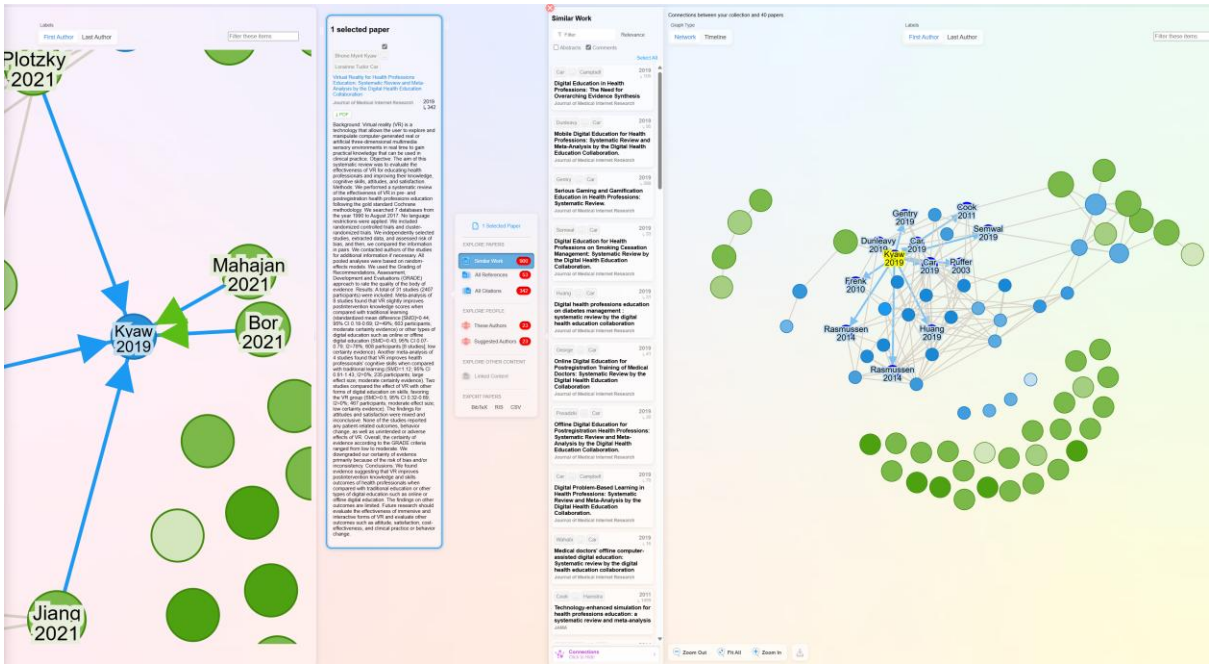

Figure 6 – Entering the "rabbit hole" of works related to the one that caught our attention

We could now backtrack to the original list of suggestions and focus on co-authorship patterns among authors of those articles. Figure 7 shows such a graph, where we see several clusters of authors with habitual co-publication within their group, but few bridges between different teams.

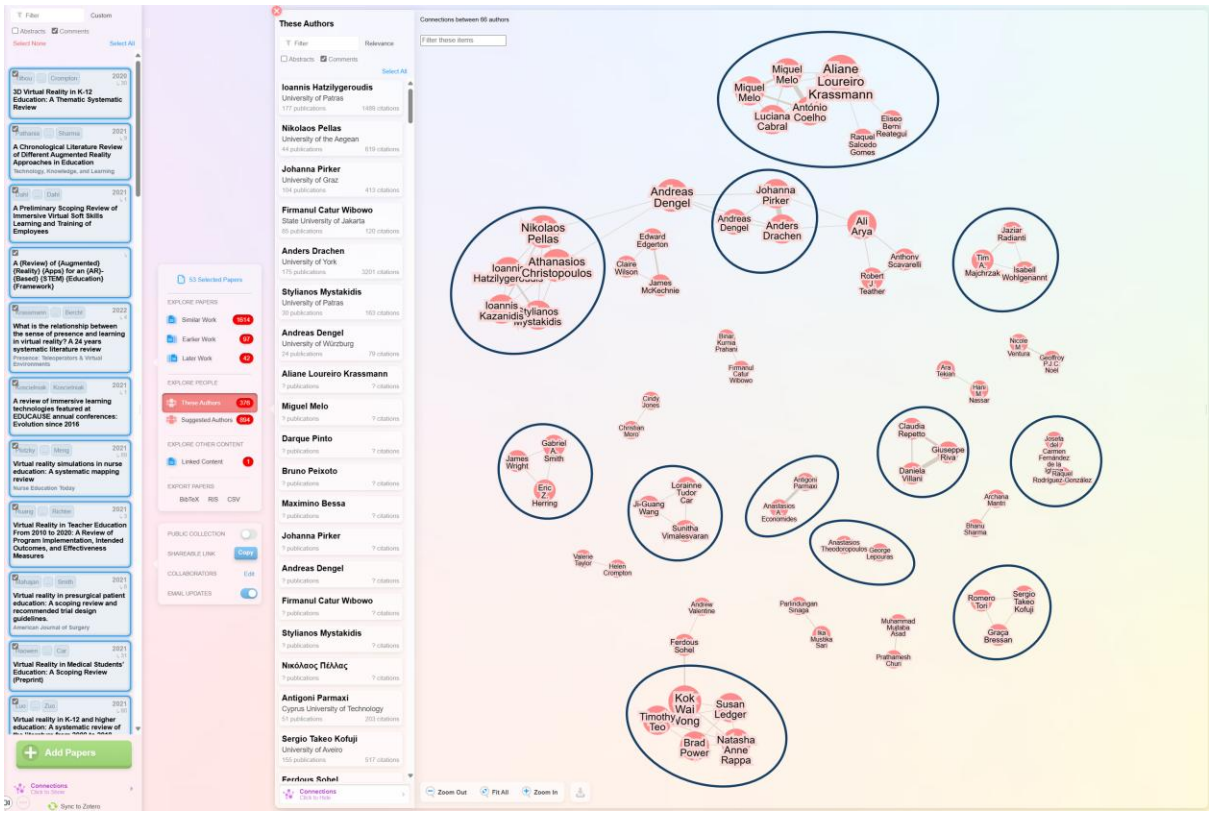

Figure 7 – Graph of authors and some of the apparent clusters

This immediately suggests that within this field there may be very different habits and practices—in terms of focus, terminology, methods and techniques, theoretical lenses, etc. If authors rarely co-author between groups, they likely remain anchored in their own group's worldview.

Being aware of this, we can explore each group's associated articles more effectively or examine the other fields these people are involved in, and who they collaborate with. This might reveal complementary perspectives we hadn't previously considered.

For example, in the **bottom cluster** of Figure 7, there's **only one co-authorship link** to someone outside that cluster: between **Kok Wai Wong** and **Ferdous Sohel**, both researchers at Australian universities. According to their Google Scholar profiles:

- Wong focuses on **computational intelligence, extended reality, and serious games**.
- Sohel specializes in **computer vision, machine learning, medical imaging, and digital agriculture**.

These themes were **not obvious** in the previous cluster map (Figure 4). If we find them relevant to our goals, we might ask: **"Who else works with Ferdous Sohel?"**, yielding the result shown in Figure 8.

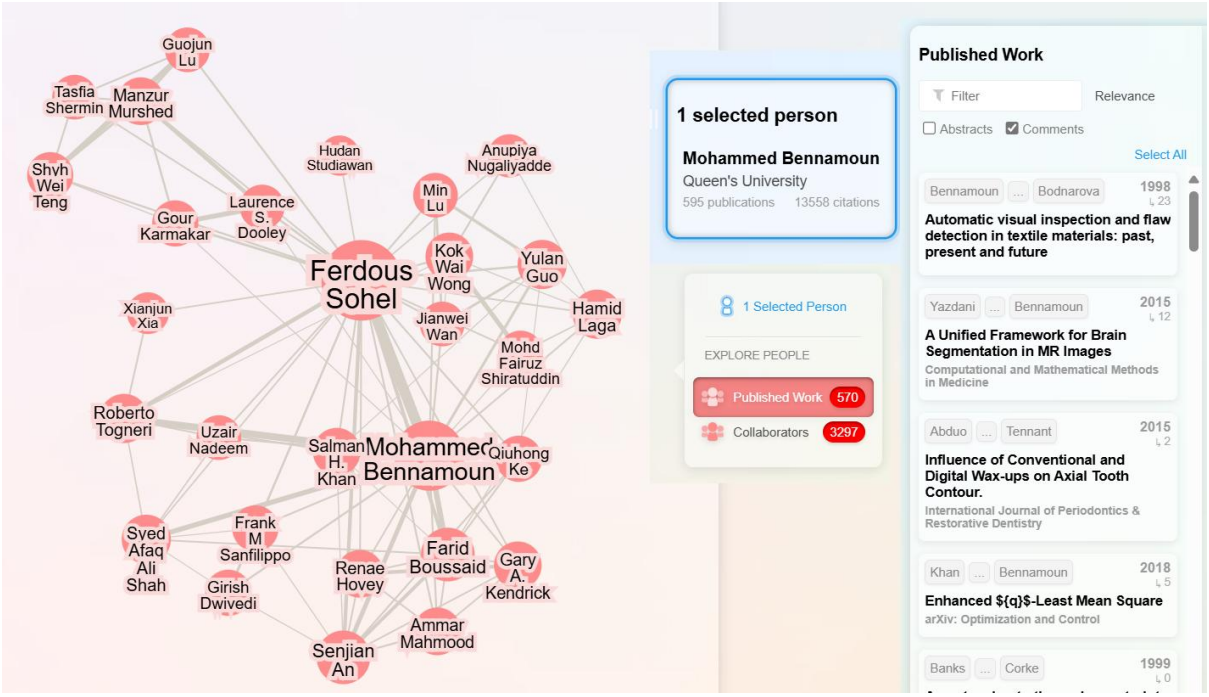

Figure 8 – Co-author network of Ferdous and articles by his most frequent co-author, Mohammed

Unsurprisingly, since Sohel's interests are diverse, his co-author network contains **several clusters**. However, there are also close and frequent co-authorships between several of these authors, indicating collaborative work. We can quickly explore their fields of study—for instance, through the work of **Mohammed Bennamoun**, Sohel's closest collaborator, which leads us into **computer vision**.

Recall that our original seeded topic was **learning with virtual, augmented, or mixed reality systems**. We're still relatively close to that topic: we can intuit a connection between **computer vision** and immersive learning (e.g., detecting aspects in images). But that connection may not have occurred to us until now: **this exploration made it emerge**, and we now know **which author networks** are working on it.

We could now explore other perspectives, such as the **citation networks** of specific articles. However, the functionality is similar, so instead, let's **leave the rabbit hole** before we get lost in it… and move on to the next tool: **Litmaps**.

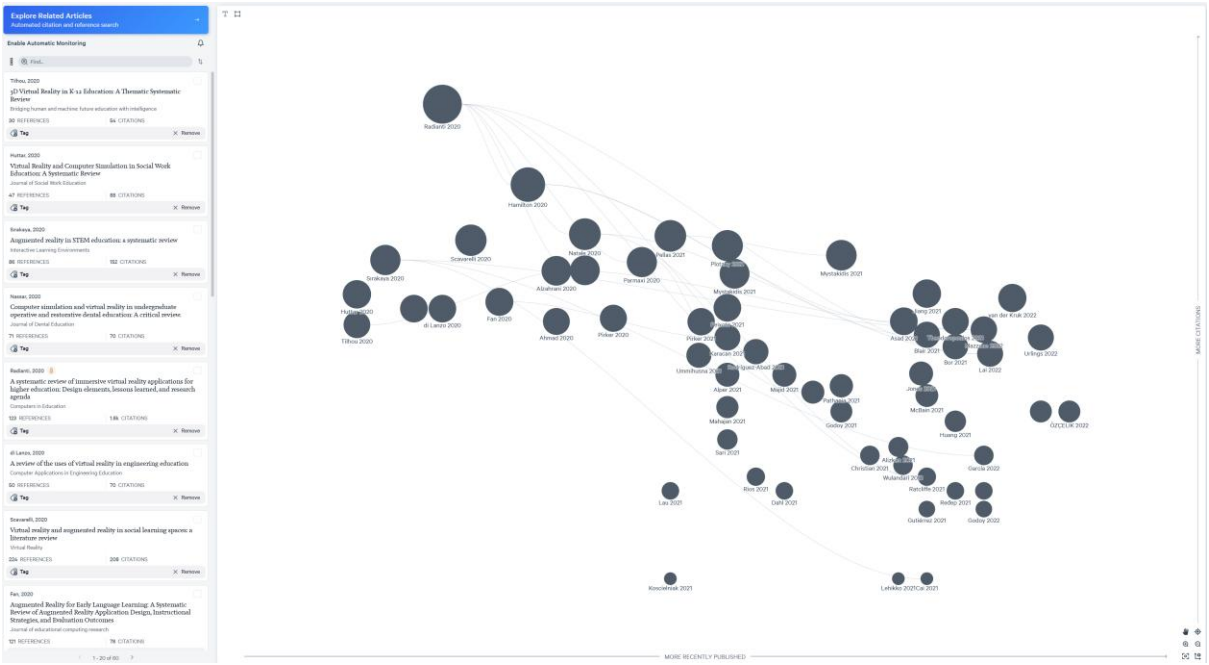

Figure 9 – Co-citation network between articles in a collection

In **Figure 9**, we see the same collection of articles on learning with virtual, augmented, and mixed reality, dated between 2020 and 2022, now visualized through **Litmaps**. The spatial layout chosen places the **oldest articles on the left** and the **most recent on the right**. The **vertical axis** represents the **number of citations**, with the most cited at the top. Note that a **logarithmic scale** is applied on the vertical axis.

Naturally, there is a **left-to-right flow**, since an article can only cite another that was published before it. But we can also observe that some earlier articles are **not only much more cited overall**, they are also **cited more frequently within this collection**.

Let's pause and reflect: a highly cited article is usually **influential globally**, but **not necessarily** influential **within our focused collection**. Let's explore this idea further.

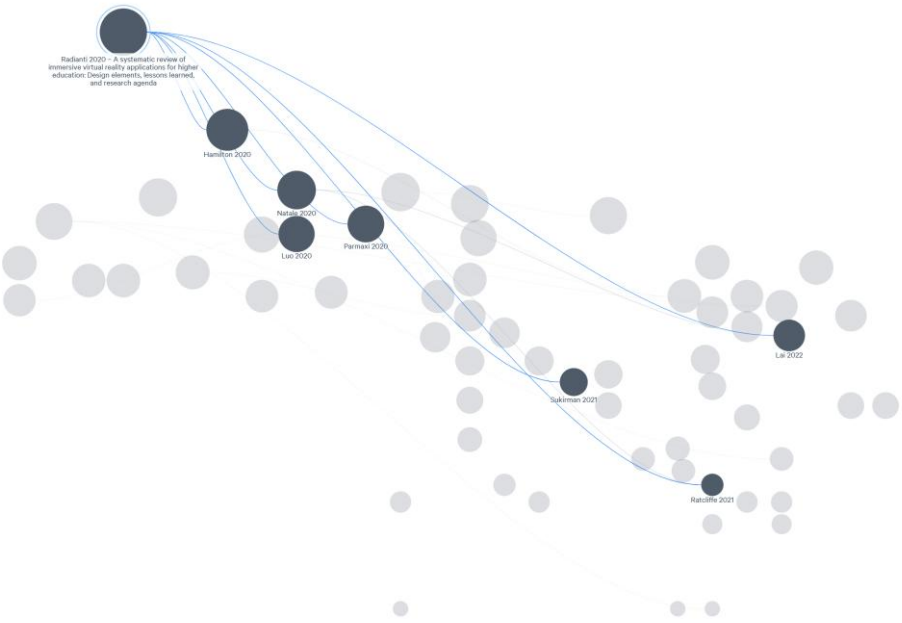

Figure 10 – The most globally cited article is also well-cited within this collection

In Figure 10, we click on the **most cited article in the collection** (>1800 citations globally). We see that it is also cited by **several other articles** in the collection. So, it is **relevant both generally and specifically**.

Now let's explore **two other articles**, also highly cited globally and within the same collection. **Figure 11** shows the contrast:

- On the **left**, an article with **224 citations**, about **virtual and augmented reality in social learning environments**.
- On the **right**, one with **152 citations**, about **augmented reality in STEM education**.

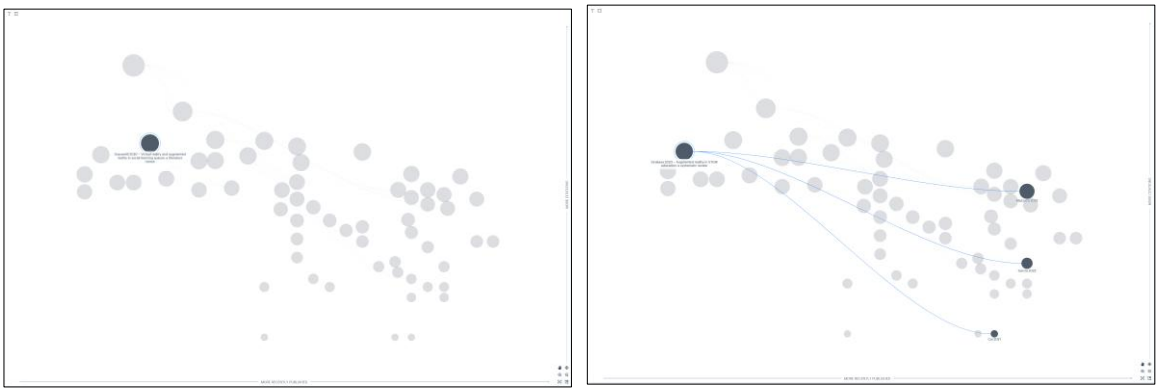

Figure 11 – Two globally cited articles with different relevance inside the collection

The **first article** (left) is **more cited overall**, and it even addresses **two technologies** (virtual and augmented reality), whereas the **second** focuses on just one. However, **within our collection**, the **first article is ignored**, while the **second influenced three later articles**.

Analysing such situations, globally, helps us **understand the nature of our collection**:

- On one hand, identifying the more influential articles allows us to prioritize them in our reading and exploration, since they likely shaped the ideas that emerged in the following years.
- On the other hand, identifying globally influential articles that had little or no impact within our collection helps us detect gaps or blind spots, especially if those articles discuss themes that we find relevant to our research objectives.

These same techniques can be applied to articles **suggested by Litmaps**, based on our collection.

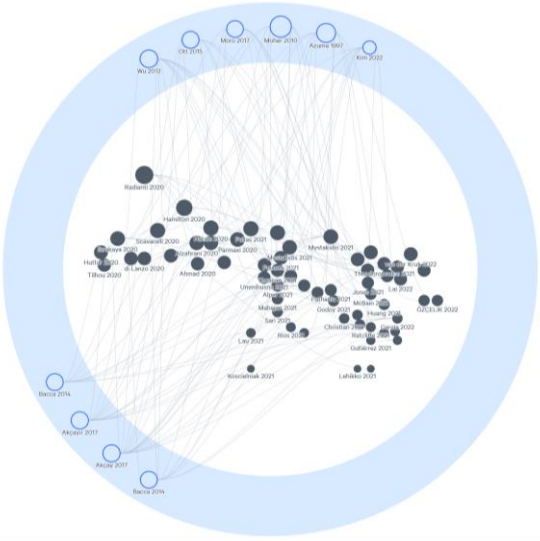

Figure 12 – Articles suggested based on our collection and their citation network

Figure 12 shows an example: when we ask Litmaps to suggest articles, it presents a diagram where the **original collection graph** is surrounded by a ring of **new articles** with **strong citation links** to those in the collection.

We can use Litmaps' filtering tools to explore these suggestions and gain **new perspectives** on our area of interest. For instance, in **Figure 13**, we applied **two types of filters**:

- On the **left**, we filtered by **publication date**, looking for articles from **2023–2025**, i.e., more recent than those in our collection (which ends in 2022). This lets us see if **previously influential articles** have remained relevant, or if **new ones have taken their place**.
- On the **right**, we filtered the suggestions for articles that included the phrase **"360 video"** in the title or abstract. This allowed us to **find works on immersive video** that weren't in the original collection and see **which of our existing articles reference them**—revealing **new perspectives, foci, or characteristics** and helping us **reassess our exploration priorities**.

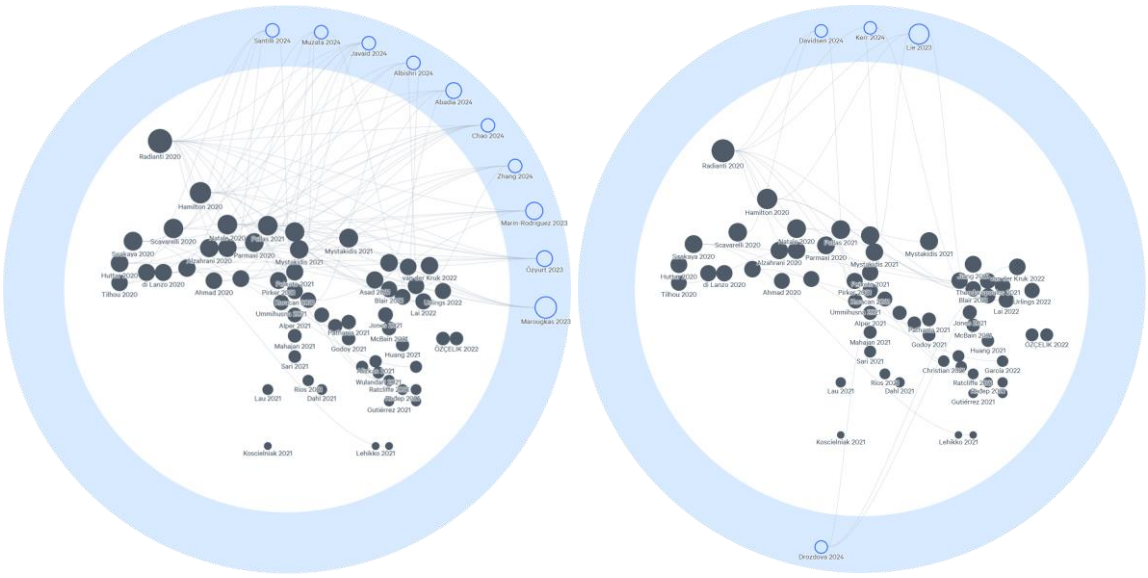

Figure 13 – Two types of filters applied to suggestions: 2023–2025 (left) and "360 video" in title or abstract (right)

Note that these tools offer many additional features. For instance:

- **Litmaps** lets you **colour-code themes** as a form of visual annotation, helping you more quickly detect connections to subtopics within the collection.
- **ResearchRabbit** allows you to explore **non-academic content** related to academic works, helping you understand how certain articles have reached **the press or influencers**.

Other tools will also offer their own exploration features.

3.5. Step 4: Analysis and Discussion of These Document Sets

After the previous steps, you will now have several **more refined sets of documents**, covering each relevant aspect of your general area of interest. You will have already **discarded some materials** you initially expected to use and **included others** that emerged as important. Even so, the path ahead is still likely to feel **unclear**, with **many branching options**.

It is now necessary, therefore, to look at the **documents and at yourself at the same time**:

- What are your **intentions**, really?
- Where do you want to **invest your effort**?
- What are your **personal goals and priorities** for conducting research?

This step remains both **technical and personal**: you must analyses the documents to identify **fundamental lines**, **contradictions**, **alternatives**, **consolidated facts**… but also to **understand yourself** as the researcher who will carry some of this work forward. You need to discover which path seems **most promising**, given your own characteristics and context.

At this stage, **various tools** can help you. That's why this step is called **"analysis and discussion"**. You can engage in **co-intelligence**, whether with teams of human colleagues, or through **mental exercises** in which you adopt different perspectives and argue with yourself, or by debating with **AI agents**.

One tool you can use here is NotebookLM. Of course, you can use other tools—this is just an example of how to leverage these possibilities.

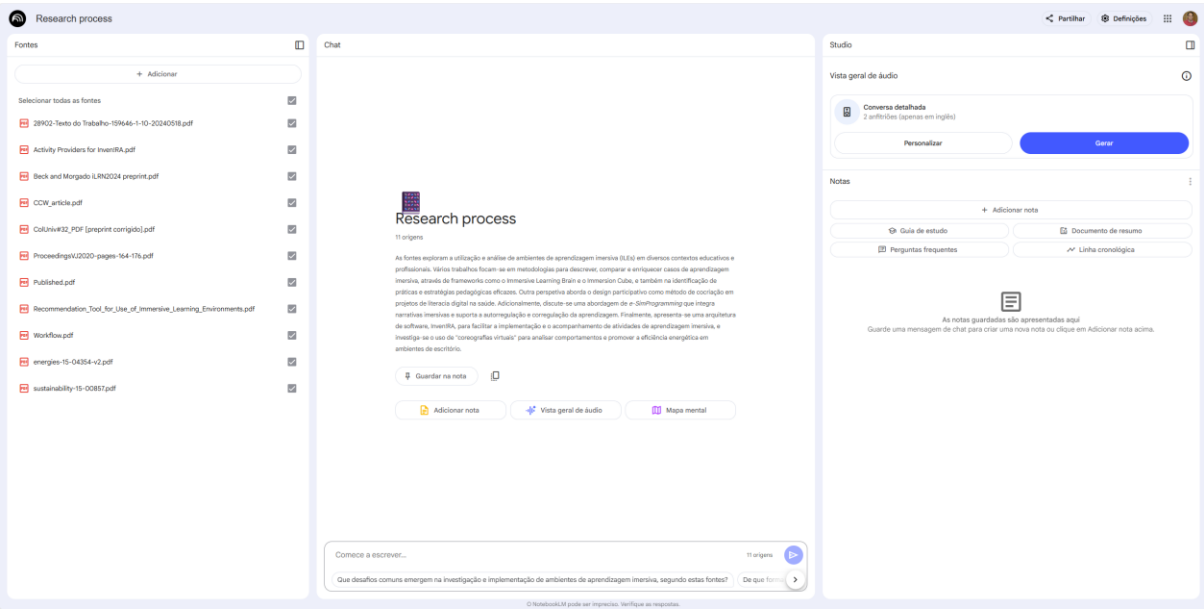

Figure 14 – NotebookLM interface during a document-based discussion

In **Figure 14**, we see an overview of NotebookLM:

- On the **left**, the loaded document set appears under "Sources".
- In the **centre**, the "Chat" is the discussion space.
- On the **right**, under "Studio", we find supporting tools.

The most immediate use is simply to **engage in dialogue**—to talk about the documents. Since you're working in **co-intelligence** with the AI, you can quickly obtain a **summary**:

> *The sources explore the use and analysis of immersive learning environments* (…) *in various educational and professional contexts. Several works focus on* (…)

You also get **suggestions for follow-up questions**, generated based on the documents:

> *– What common challenges emerge in researching and implementing immersive learning environments, according to these sources?*
> *– How do different pedagogical and technological approaches relate to the context of immersive learning discussed?*
> *– What are the main trends and models for analyzing and designing immersive learning experiences presented?*

These three prompts already highlight **three different perspectives** on the same document set:

1. The **challenges** in research and implementation;
2. The **relationship between pedagogy and context**;
3. **How to analyse and design** learning situations.

Your **intellectual posture** at this stage should be one of **awareness**: awareness of this diversity of perspectives… and a willingness to **deepen** your understanding of them.

To support this process of awareness, you can go beyond just the discussion and explore other NotebookLM features, found in the **Studio** panel. The **intellectual approach remains the same**: explore to gain awareness.

For example, you can request a **mind map** of the documents. **Figure 15** shows one such map generated automatically. Many items include **“>” and “<” symbols** to help you expand and collapse concept branches.

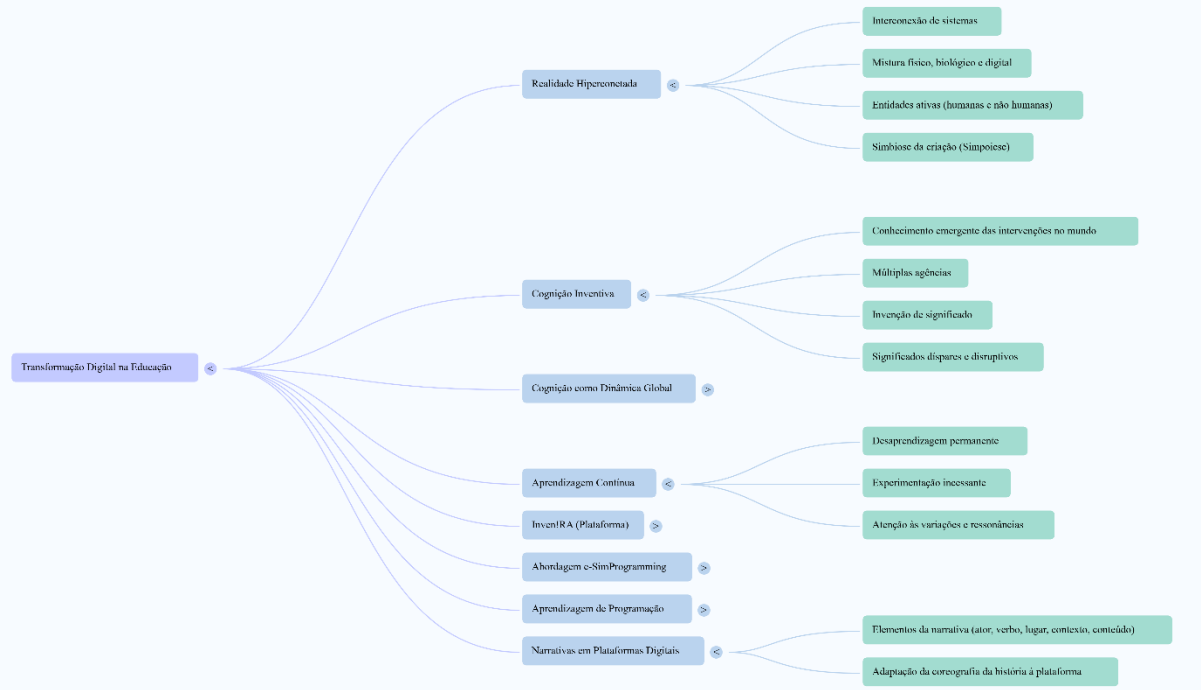

Figure 15 – Mind map produced by NotebookLM based on the supplied documents

Another tool is the **audio overview**, which consists of the automatic creation of a **podcast episode**, where **two speakers debate** the themes of the articles—see **Figure 16**.

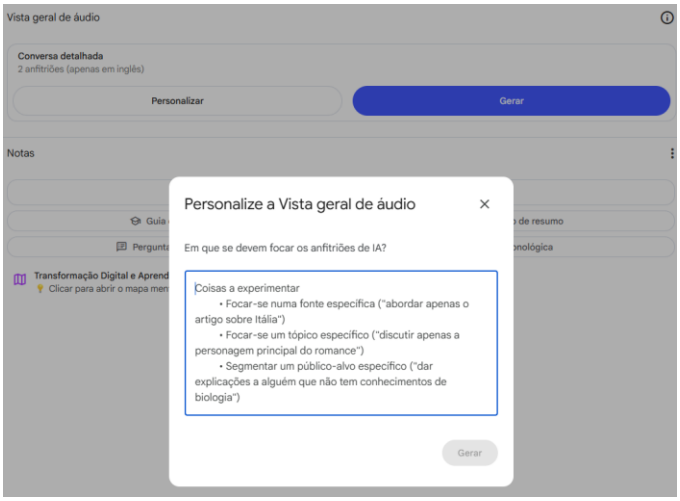

Figure 16 – Audio overview setup

This means you can **go for a walk**, **do your shopping**, or carry out household tasks… while **listening to a lively debate** between two hosts about the themes you've been studying. Or several debates! You can generate **multiple podcast episodes** by using different instructions.

This **diversity of ways** to explore a document set should lead you to **take notes**—perhaps quick notes in NotebookLM's Studio section, but as your ideas solidify, you should register them in your **personal knowledge system** (like Obsidian). Of course, you can also **bring previous notes from Obsidian** into NotebookLM to inform the discussion.

All other features—like "Study Guide", "Summary Document", "FAQs", "Timeline", etc.—serve the same purpose: **To help you explore and gain awareness of your field**, your possibilities, your interests and goals, and the divergences, convergences, contradictions, and convictions within it.

3.6. Step 5: Reflection to Synthesize the Topic into Concept, Context, and Perspective

During and after the awareness process in the previous step, it is important not to lose sight of this exercise's primary purpose: choosing your research path.

If we let this reflection remain in the abstract realm of ideals, we may easily fall into an unrealistic or misguided process. We must confront our concepts with reality, as Seymour Papert urged us:

> "(…) *subject to the test of reality; if they don't work, it becomes a challenge to understand why and overcome the obstacles. They can be shown, shared, and debated with others.* (…) *they become transitional objects for the personal appropriation of ideas."*
>
> (Papert, 2019)

So, in this exercise, we're discovering our **research intent**—but it must be made concrete. A major risk here is that we **hide behind eloquent but hollow phrases**, full of abstractions and generalities.

To ensure that our idea, once turned into an intent, **passes the test of reality**, the proposed approach is to **write it using a three-part structure**: **concept**, **context**, and **perspective**.

- The **concept** is the central idea. It may be the area of interest or the theme you want to explore.
- The **context** is the setting in which you want to analyse the concept. Where does the concept occur, where is it applied, who is involved, etc.—these are ways to frame the concept.
- The **perspective** is how you want to approach the concept within that context.

You can make this even more concrete by writing a **descriptive sentence** explaining how these three parts come together. This helps uncover contradictions: e.g., if the sentence introduces concepts, contexts, or perspectives that weren't in the original formulation.

Let's now look at **some examples** to support your reflection.

3.6.1. Examples of Intents in the Format: Concept | Context | Perspective

`[concept:]` Thematic networks `[context:]` of science `[perspective:]` visualization methods

> This means the aim is to find **methods for visualizing thematic networks**, where those networks pertain specifically to **scientific domains** (not other types of themes).
> The output will be sets of **methods**—types of diagrams, matrices, and so on.

In other words: the study focuses on the **concept** of thematic networks in general, pays attention to their **application in science**, and seeks to discover **what visualization methods** are used—that's the chosen **perspective**.

`[concept:]` Visualization methods `[context:]` of thematic networks `[perspective:]` for science analysis

Here, the focus is on **ways of analysing science** using **visualization methods**, particularly those that visualize **thematic networks**.
The goal is to find **ways to analyse science**: identifying relevant themes in networks, distinguishing peripheral and central themes, etc.
So, the concept (visualization methods) is treated generally, the context is thematic networks, and the perspective is on **what kinds of science analysis** these visualizations enable.

`[concept:]` Instructional design models `[context:]` for immersive virtual reality environments in online higher education `[perspective:]` with their phases and characteristics

This means the intent is to identify the **phases and characteristics** of instructional design models specifically created for **immersive virtual reality environments** used in **online higher education**.
The results will be these phases and characteristics, as long as they apply in that context.
In other words: the study considers instructional design models in general, focuses particularly on those for immersive VR in online higher education, and investigates the **phases and features** of those models—that's the perspective.

3.7. Step 6: Iteration over Previous Steps to Refine or Redirect the Focus

The result of your analysis, discussion, and reflection may be the realization that you need to explore a topic more deeply, seeking out new documents. In most cases, the smartest move is to go back just one step, to where you explored the most influential and foundational documents on that topic.

However, you should always keep the freedom to go back further, if you truly consider it relevant. That said, be mindful and reflect carefully whether such a move is due to necessity… or just anxiety.

4. Second Stage: Refinement and Focus

4.1. Objective: Delimit and Structure the Topic

Once the first stage is complete, your scope is defined: you have a clear intent, expressed through a concept, a context, and a perspective.

Now, you need to find the evidence to support that intent. To do so, you don't want to get lost in a long narrative full of endless details. You want to establish a conceptual structure to help you interpret what exists, navigate through it… and eventually, build upon it.

That's why we also follow a structured process here, supported by AI tools, but with a different goal (see Figure 17).

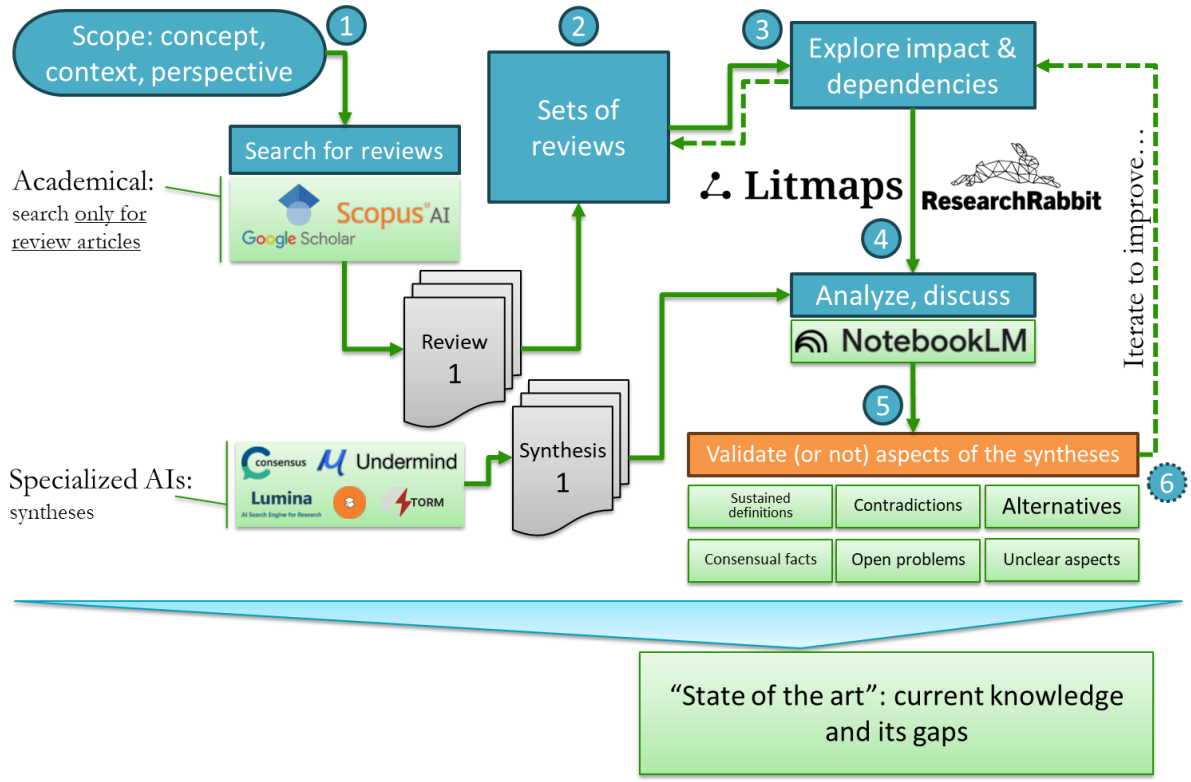

Figure 17 – Process for the second stage: finding the evidence

In this new process, the goal is no longer to **discover your intent**, but to **structure the existing evidence** surrounding that intent—in other words, to understand the **state of the art**.

Of course, building a structure always involves a **degree of subjectivity**. When organizing current knowledge, we'll **prioritize certain criteria**, apply **preferences**… These decisions should be as **conscious as possible**, so that later we can **revisit** or **challenge** them, and reduce the illusion of having built the "right" structure.

4.2. First Step: Prepare the Confrontation Between Reviews and Syntheses

With your intent clearly defined (**concept, context, perspective**), you're now ready to prepare two "opponents" for a **conceptual debate** that will help structure the area:

- A collection of **literature reviews** on the topic/intent.
- A collection of **syntheses** about the topic/intent produced by **specialized AI tools**.

The key difference between this step and the one in the first stage is that now you must be **more systematic**:

1. In the search for scientific articles, you will deliberately focus on panoramic works—that is, literature reviews. You may find many or few; if too many, you might have to narrow the topic, limit your analysis to more recent works, or assess methodological quality to focus on the most rigorous ones.
2. When gathering AI-generated syntheses, you will now use specialized systems focused on literature reviews and apply them explicitly to your specific intent (whereas previously, you were still operating at a more general, exploratory level).

4.3. Second Step: Document Sets

As in the previous stage, the outcome here will be **separate sets of documents**: one set composed of the **literature reviews** you selected; another set composed of **syntheses** produced by the **specialized AI tools**.

4.4. Third Step: Exploring the Impact and Dependencies of the Reviews

With your set of literature review articles (remember that at this point you are only working with reviews focused on your intent), you can now reuse the tools for exploring citations and conceptual proximity—such as ResearchRabbit and Litmaps.

The goal remains the same as before: to identify influential, foundational articles that may have escaped your attention, but that are cited by several of the reviews you selected.

Note that in this step, you are not yet using the AI-generated syntheses, because those are not actual articles—they are summaries created on the spot by intelligent systems.

4.5. Fourth Step: Analysis and Discussion Through Confrontation

After the previous step, you will have a refined set of literature reviews, which you will now confront with the syntheses produced by the AI tools.

Once again, you'll use support tools such as NotebookLM to analyse these two document sets, both individually and through comparison and contrast.

4.6. Fifth Step: Validating (or Not) Aspects of the Syntheses

The analysis and discussion processes should now begin to **converge** toward a **more critical and reflective understanding** of the field. This is where the **conceptual structure** you're aiming to build will take shape.

As your analysis and discussion progress, you should begin extracting notes that will allow you to establish:

- the **current state of knowledge** in the field,
- an **awareness of the evidence**,
- in short: the **state of the art**.

To achieve this, you should be able to identify and articulate, from the confrontation between **AI syntheses** and **literature reviews**, the following:

- **Supported Definitions:** What definitions are given for the key terms and concepts? What are the original sources of those definitions? Note: It's likely there will be more than one definition per concept or term—hence the plural: **definitions**. Developing your understanding of the field means recognizing the variety of perspectives even at the definitional level.
- **Contradictions:** Contradictory viewpoints are common in all fields. These may involve conflicting definitions; opposing assessments (e.g., "X supports Y" vs. "X harms Y"); or other disagreements. Such contradictions must be noted and recorded so that you are aware of their existence. It is too early to evaluate the merit of each side—that may require deeper

conceptual understanding or further analysis. For now, what matters is recognizing the contradiction and its nature.

- **Alternatives:** Not all divergences are contradictions. Often, they are **valid alternatives** you need to be aware of, so you don't **misinterpret** findings. For example: for some authors, *"immersion"* refers to being surrounded by synthetic spaces (e.g., "immersion in virtual reality"). For others (me included), it means the **mental state of absorption**, losing awareness of physical surroundings—something that can happen with virtual reality, but also while reading, watching a film, or getting deeply involved in a task or thought. Still others define it as **being and living within the target reality**—for example, learning a language by living in a country where it is natively spoken.
- **Consensual Facts.** In addition to divergences, you also need to recognize what is generally accepted as fact. This doesn't imply total unanimity—there are dissenting voices in nearly every area—but some facts are so widely accepted that questioning them is considered bizarre or implausible by most experts. These are the assumed foundations that don't require lengthy explanations in most contexts.
- **Open Problems.** Another key element is identifying the problems that **specialists themselves** acknowledge as **unresolved** or **poorly addressed**: Due to uncertainty in the available data; Due to a **lack of data** to validate or invalidate hypotheses; Or due to **conflicting findings**.
- **Unclear Aspects.** Finally, be on the lookout for elements that lack **clarity**—whether in definitions, in factual claims, or in the framing of problems. Examples:
    - A factor *X* is generally considered to increase an effect *Y*, but the **degree, speed, or conditions** of this relationship are unclear.
    - A concept is often used **intuitively**, but there are **no objective ways** to measure or compare it.
    - People attribute causes or effects to a **complex concept** (like *learning*), but do so **without explanation or caution**, as if it were an objective, self-evident fact.

  These are all aspects that should be **flagged as lacking clarity** in the field.

4.7. Expected Outcomes: Initial Structuring of the Topic and Identification of Gaps

By the end of the previous steps, you will have more than just a list of concepts—you will have an **organized structure**:

- of the **concepts themselves**,
- of their **interrelations**, **conditions**, and **uncertainties**,
- and also, of their **omissions**.

This structured and evaluated web of concepts is your goal: the **state of the art**.

5. Third Stage: Judging the Need for an Update via Systematic Review

5.1. When Do We Stop?

After comparing the **literature reviews** and the **AI-generated syntheses**, both focused on your intent (expressed through **concept, context, and perspective**), you will now have a **valuable understanding** of the knowledge available on your topic—i.e., the **state of the art**.

But we know the topic is likely **not exhausted**.

Even if you continue exploring it for **years or decades**, there will almost certainly be **something you've missed**:

- An author using **unexpected terminology**,
- A study published in **another language**,
- A concept framed from a **different disciplinary angle**…

Sounds unlikely? But it's not. There are many well-known examples:

- The work of Belarusian psychologist **Lev Vygotsky** remained virtually unknown in the West for decades after his death (1934). It was only translated from Russian into English in 1962 and became widely recognized in the 1970s–1980s. Today, he is a cornerstone of educational psychology and pedagogy.
- Similarly, **Gregor Mendel's** 1866 paper on pea experiments—the foundation of modern genetics—went virtually unnoticed for more than 30 years. Why? It was published in a regional journal, in German, using mathematical and statistical language unfamiliar to naturalists at the time.

The same could be happening today—and very likely is—given the much **greater volume of scientists** and **scientific production** in our era.

So, the decision to stop reviewing the state of the art **is not based on knowing everything**. We stop when we **weigh the risks and benefits** in terms of our **current and future efforts**.

Dear student… You're about to invest many hours, starting soon, in hands-on work. Depending on your research topic and method, this time investment may also come with **financial or logistical costs**.

You want to proceed **in an informed way**, so as not to waste that effort. That means avoiding:

- **"Reinventing the wheel"**: doing something that seems original but is actually a repetition of existing work.
- Presenting claims that amount to **"I think this is good, trust me"**. (Of course, it's never phrased that way—it's dressed up in elaborate rhetoric, masking fragile reasoning.)
- Falling into **improvised analyses**: doing a lot of data collection and technical work, only to realize later you **don't know how to interpret it**, **compare it**, or **explain what it means**.

These disasters are common.

Reinventing the wheel can happen even when someone **knows the field very well**—because **another person** might have published similar work **at the same time**.

But if we really understand the **state of the art**, we can assess whether that similar work:

- stemmed from the same assumptions,

- used the same evaluation criteria,
- followed the same principles.

It's unlikely that all those factors will coincide. And even if they do, having that **awareness** gives us the ability to make decisions on **how to meaningfully build on** what's already been done—**preserving the value of our work**.

Those who lack that awareness tend to panic and activate the **"complicator"**: inventing a highly specific, rare case just to say, *"but mine, ladies and gentlemen, is different!"* Yes, different… but likely **irrelevant**.

Wouldn't it be a better investment of time and energy to **fully understand the state of the art** so that, when surprises come up, you can **choose the most relevant direction forward**?

That's the **core decision at the end of the second stage**: if you feel you now have **adequate awareness** of your intent and that the **risk-benefit balance** is manageable, then you **stop here**.

If you feel the horizon is still **clouded**, then you **move forward**.

Often, this will be the case. The advantage is that you now have a **strong grasp of the state of the art**, which allows you to take the next step in a much **more focused way**.

Instead of gathering reviews and syntheses or conducting informal searches, you can now proceed to a **systematic literature review**—one that includes **original research articles**, not just reviews.

Thanks to your solid foundation, you'll be able to:

- define **relevant search terms**,
- set up good **inclusion/exclusion criteria**,
- and establish **clear analysis goals**.

There are several protocols for this. Personally, I follow those by **Barbara Kitchenham (2016)**—but they are not the only ones.

There are also tools to help you with this process, such as Parsif.al.

5.2. How Can We Tell if We Have Reached Adequate Awareness?

A key recommendation for any beginning researcher is to **seek advice from your supervisor(s)**. The experience of the supervisory team is usually a very good guide at this stage.

However… we're not always adequately supervised. That's just reality.
And even when we are, supervisors don't always respond **promptly** or with enough **attention**.

An alternative is to **present your ideas at a scientific event** in your field. Share them there, and then take full advantage of the **informal conversations** with experts attending the event, to check:

- Are you overlooking **essential issues**?
- Are there **terminological problems** you're unaware of?
- Are you citing **facts already disproven**?

This is also an excellent way to **break out of the tunnel vision** that often develops within research teams—exposing yourself to **different theoretical lenses** and **alternative conceptions** of the problem.

Many events offer specific sessions for early-stage researchers, often called **"doctoral consortiums"** or similar, designed to facilitate this kind of exchange. But…

- High-quality scientific events are **not held every week**, nor just around the corner.
- **Registration fees** can be expensive.

So, it's important to have **some way** of building awareness and confidence that you are **theoretically saturating** your field (yes, that's the jargon used).

Here are two particularly **rules of thumb**:

1. **Symptoms of Completeness.** You start noticing that the contributions of various articles are repeating—they start to feel "like more of the same." Your new notes begin to interconnect with your previous ones in ways that feel familiar and stop surprising you. The foundational references in the articles start to become repetitive…
2. **A Sense of a Rich Overview.** You begin to feel that you have a broad and diverse view of how different researchers approach the field: different theories, different methods… Not just because you rushed to make a list, but because you insisted—you spent time, and the surprises faded, while a kind of *déjà vu* set in.

Figure 18 gives an example of how this can manifest. In that study, my co-author and I already had a broad sense of the different perspectives on the concept of "immersion", thanks to a prior review of the state of the art. This allowed us to identify which perspective each author at a conference adopted regarding the concept—and to identify which theoretical references supported each of those views. The resulting diversity was unsurprising—it confirmed our earlier understanding, but now with a rich overview that gave us more confidence to assess the risk-benefit relationship.

Immersion perspective	***No. papers[a]***	***Main ref. on immersion[b]***	***No. papers[a]***
Technological fidelity	7	Slater (2003) or Slater & Wilbur (1997) or Azuma (1997)	3
Unspecific	4	None	6
Engagement	2	Csikszentmihalyi (1992) and/or Kolb (1984)	1
Interaction	1	Barab (2007)	1
Narrative	1	Whitton (2011)	1

a. Sum of immersion perspective is not n=12 because three papers employ two perspectives.

Figure 18 – Example of a theoretical overview, from Morgado & Beck (2020)

6. Conclusion and Final Reflections

This process of exploring your area and your intent is the **foundational part of your research journey**. If done with **passion and ambition**, it is **transformative** and **empowering**. More than enabling you to **speak persuasively**, it allows you to **decide confidently and clearly**. It gives you a **compass** to help you choose what to explore, where to go, and how. Later, you will **naturally and regularly repeat** this process—throughout your research and throughout your career. Sometimes more hastily, other times more ambitiously. It becomes a **natural part of scientific competence**.

In other words, it is an **iterative process**—not something you do to **finish**, but to **get started**… well launched, but on a **permanent journey**

There's one last detail: **What about sources beyond academic publications?** Aren't they also relevant? Things like:

- **Code repositories** like GitHub, where most content hasn't been reviewed or analysed in academic articles;
- **Corporate newsletters**, where early glimpses of innovative techniques are shared;
- **Social media posts**, where people report something new that just worked, something that might not be documented anywhere else…

Are we ignoring the **real** state of the art?

No, we are not! Those sources are **field data**, not scientific knowledge!

Just as a **biologist** studies plants and animals, an **information scientist** can study **code**, **technical sources**, or **user reports**. But the knowledge isn't in the "plants" and "animals" themselves.

Code, technical documentation, reports—they're our "flora and fauna." We can and should analyse the field—but **only after we have the theoretical lenses** to do so, built from the **state of the art**.

Or the reverse: we can let what we see "in the wild" inspire us and then build lenses to analyse it.

7. References

Copeland, S. (2019). On serendipity in science: Discovery at the intersection of chance and wisdom. *Synthese*, *196*(6), 2385–2406. https://doi.org/10.1007/s11229-017-1544-3

Gerasimov, I., KC, B., Mehrabian, A., Acker, J., & McGuire, M. P. (2024). Comparison of datasets citation coverage in Google Scholar, Web of Science, Scopus, Crossref, and DataCite. *Scientometrics*, *129*(7), 3681–3704. https://doi.org/10.1007/s11192-024-05073-5

Kitchenham, B. A., Budgen, D., & Brereton, P. (2016). *Evidence-based software engineering and systematic reviews*. CRC Press.

Martín-Martín, A., Thelwall, M., Orduna-Malea, E., & Delgado López-Cózar, E. (2021). Google Scholar, Microsoft Academic, Scopus, Dimensions, Web of Science, and OpenCitations' COCI: A multidisciplinary comparison of coverage via citations. *Scientometrics*, *126*(1), 871–906. https://doi.org/10.1007/s11192-020-03690-4

Morgado, L., & Beck, D. (2020). Unifying Protocols for Conducting Systematic Scoping Reviews with Application to Immersive Learning Research. In *2020 6th International Conference of the Immersive Learning Research Network …*.

Papert, S. (1999). Introduction: What is Logo? And Who Needs It?" In *Logo philosophy and implementation*. LCSI. Available online: http://www.microworlds.com/support/logo-philosophy-papert.html