

Management Innovation Lab @ HOST



Working Paper 01/26

AI in Foresight: AI-supported Trend Radars

Design | Implementation | Maintenance

Author

Prof. Dr. Björn P. Jacobsen | Stralsund | Germany

Copyright Information

Björn P. Jacobsen. All rights reserved.

Published by

Management Innovation Lab @ HOST
<https://www.hochschule-stralsund.de/mil-1/>

Edition

1.0. September 2026

Disclaimer

Although the author has made every effort to ensure that the information in this publication was correct when it went to print, the author does not assume and, with this, disclaim any liability to any party for any loss, damage, or disruption caused by errors or omissions, whether such errors or omissions result from negligence, accident, or any other cause.

ISSN Information

-

Connect

bjoern.jacobsen@hochschule-stralsund.de | <https://www.hochschule-stralsund.de/jacobsen/>

Cover Picture

perplexity.ai

General Disclosure

Perplexity and *FIBRES* were used in the assessment reported in this working paper. The developer of *FIBRES* provided free access for the purposes of this research.

AI Disclosure

Artificial intelligence tools were used to prepare this working paper and the exploratory test reported in Section 4.3. *Perplexity* supported source discovery, research planning, evidence organisation, preliminary analysis, drafting, and the production of the first test's radar and associated materials. *FIBRES* was used to configure and automatically populate the second test's draft radar, including its linked items, network visualisation, and report. Section 4.3 describes these uses, their outputs, and the extent of human review. *Grammarly* assisted with language editing, and *Perplexity* created the cover image.

AI-generated material was treated as assessment material, not as independent evidence or an automatically approved finding. The first test included documented human review gates and produced an approved first-draft radar; its management summary had not received final content and release approval at the time documented. The *FIBRES* radar remained an automatically populated draft without documented post-generation human approval. The author is responsible for the paper's research design, interpretation, conclusions, and final text, including any errors or omissions.

Table of Contents

1.	Management Summary	4
2.	Introduction	5
2.1	Trend Radars in Innovation Management	5
2.2	Challenges in the Application of Trend Radars	6
2.3	Purpose and Scope of the Working Paper	7
2.4	Research Design and Paper Structure	7
3.	Theoretical Framework	9
3.1	Foresight Process in Innovation Management	9
3.1.1	Innovation Management and Future-Oriented Intelligence	9
3.1.2	Foresight, Uncertainty and Discontinuity	10
3.1.3	Horizon Scanning and Weak Signals	10
3.2	Trend Radars in Foresight Work: Purpose, Functions and Architecture	11
3.3	Trend Radar Process	12
3.3.1	Source Selection and Horizon Scanning	12
3.3.2	Signal Extraction and Coding	13
3.3.3	Clustering and Trend Formation	13
3.3.4	Evaluation, Prioritisation and Scenario Interpretation	13
3.4	Conventional Trend Radar Process and Data Practices	14
3.4.1	Define Search Areas	14
3.4.2	Define Source Strategy	15
3.4.3	Identify Signals	15
3.4.4	Categorise Signals	16
3.4.5	Build Trend Radar	16
3.4.6	Create Narratives	16
3.4.7	Disseminate Results	17
4.	AI-supported Trend Radar Approaches	18
4.1	Classification of relevant of AI Tools	18
4.2	Market Scan of AI-Supported Trend Radar Tools	19
4.3	Test Case: “I Feel Good”	20
4.3.1	Test 1: Analyst-led Workflow Using Perplexity	23
4.3.2	Test 2: Automated Population Using FIBRES	24
4.3.3	Cross-case Assessment	25
5.	Recommendations for Implementation	27
5.1	Define Purpose, Decision Context, and Governance	27
5.2	Build the Data and Source Foundation	27
5.3	Combine AI Automation with Expert Review	28
5.4	Maintain a Living Trend Radar	28
5.5	Design for Visualisation, Communication, and Decision Integration	29
6.	Conclusion	30
7.	References	31
	Annex	33

1. Management Summary

Trend radars can help organisations identify developments relevant to future innovation. Their value depends on credible sources, careful interpretation, regular updating, and clear responsibility for decisions. This working paper examines how artificial intelligence (AI) can support trend radar work for food processing small and medium-sized enterprises (SMEs) while preserving human judgement.

The working paper combines a conceptual framework, a market scan, and two exploratory tests using the same research brief on nutrition, beauty, and health. It investigates which stages AI can support and how SMEs can use it without forgoing accountability. The tests compare an analyst-led *Perplexity* workflow with an automated radar population in *FIBRES*. They do not rank the tools or assess predictive accuracy.

AI can assist with source discovery, summarisation, signal extraction, clustering, trend drafting, visualisation, and reporting.

The *Perplexity* test organised these tasks through review gates. After analyst-led expansion of its source pool, the register contained 132 records; review resulted in 38 approved signals, 11 clusters, and ten positioned trends. The *Perplexity* output disclosed evidence gaps, while its management summary remained pending final approval.

FIBRES demonstrated rapid structuring and visualisation. It synthesised 58 phenomena and produced a populated radar, network view, report, and exportable records. Some material, however, had an unclear connection to the specified food product decisions. The record does not document completed human approval of sources or radar positions; its output is an automated draft for review.

Neither extensive retrieval nor a populated radar establishes reliable strategic insight.

Perplexity provided a clearer documented review trail but required substantial analyst effort. *FIBRES* shortened the path to a visual draft but required relevance filtering and evidence appraisal. The tests did not establish comparable human effort, costs, or performance across repeated update cycles.

SMEs should define the radar's decision purpose and approval roles, maintain a source list/register, review AI-generated material before inclusion, and document why signals become trends.

The objective of the trend radar is a living evidence base for better-informed innovation decisions - not an automatically completed picture of the future.

2. Introduction

Organisations need to identify developments that may affect future innovation before they are reflected in market or performance data. Trend radars can support this task by structuring and communicating future-oriented information, but conventional trend radar work is resource-intensive and difficult to maintain, especially for SMEs. This working paper assesses **how AI can support selected stages of trend radar** design and maintenance while retaining transparent evidence and human judgement. It develops this assessment through a literature review, a market scan of tools, and an exploratory test using a food industry case.

2.1 Innovation Management & Trend Radars

Innovation management concerns the identification, development, selection and implementation of new or improved products, services, processes, technologies and business models. It starts not only with current **customer demand** but also with the possibility that changing **external conditions** may create new opportunities or threats.




Figure-1: Trend Radars in the Innovation Management Framework (Jacobsen, 2025)

This challenge is relevant in food processing. Food businesses operate at the intersection of agriculture, processing technologies, packaging, distribution, retail, and consumer behaviour. Food innovation must therefore reconcile competing requirements, including sensory quality, safety, affordability, nutritional value, sustainability, technical feasibility, supply security, regulatory compliance, and consumer acceptance. Consumer expectations regarding health, convenience, sustainability, and plant-based products increase the need for forward-looking innovation (Asioli et al., 2017; Bigliardi and Galati, 2013).

For managers, **relevant questions** include:

- Which emerging needs, technologies, channels or regulations may affect future customer value?
- Which developments could create product, process or business-model opportunities?
- Which changes may threaten current capabilities, supply systems or market positions?
- Which areas (topics) justify monitoring, experimentation or capability building?

Trend radars can provide a practical interface between external intelligence and such management questions. They make future-oriented information visible, structured, and discussable across management, product development, R&D, marketing, procurement, sustainability, and regulatory affairs.

For SMEs, trend radars matter because they let them focus limited resources on developments most relevant to innovation. Their value, however, depends on whether they can be developed and maintained without creating a workload or technical complexity that exceeds the SME's capacity.

2.2 Challenges in the Application of Trend Radars

Trend radars can support innovation management, but their practical application is demanding. The difficulty is not the visualisation, but the **ongoing work** of identifying, selecting, validating, interpreting, updating and communicating relevant information. For the purpose of this paper, **five typical challenges** are distinguished. They provide the analytical basis for the test in Section 4.3:

Time-to-Insight Challenge - Conventional radar work can take weeks. Teams must agree on a research scope, identify and assess sources, collect information, develop categories, extract potential signals, and describe trends. By the time the first trend radar is complete, the evidence base or decision context may have shifted. The objective is therefore not fully automated strategic judgement, but a transparent, source-linked first-draft trend radar within one working day of the research brief's creation.

Signal-to-Meaning Challenge - Horizon scanning generates large volumes of heterogeneous material, including scientific publications, patents, product launches, regulatory documents, trade reports, start-up activity and consumer signals. The challenge is to transform this material into meaningful, evidence-based signals, clusters, and trends.

Especially weak signals are fragmented, ambiguous and uncertain; their relevance cannot be derived from volume or visibility alone (Ansoff, 1975; Hiltunen, 2008; Saritas and Smith, 2011). Organisations must distinguish evidence from publicity, recurring patterns from temporary attention, and strategically relevant developments from general "noise".

SME-Capability Challenge - Large organisations may maintain dedicated foresight units, trend scouts, analysts and external research. Most SMEs cannot. Their managers combine innovation work with daily operational responsibilities. Resource constraints, short planning horizons and limited methodological expertise are barriers to SME foresight (Gershman, Bredikhin and Vishnevskiy, 2016). A suitable approach must therefore be proportionate, understandable, repeatable and feasible without extensive technical infrastructure.

Living-Radar Challenge - Many trend radars remain static, one-off outputs prepared for workshops or reports. Their relevance declines fast if they do not reflect new evidence, fading signals, changing regulations, or emerging competitors. A useful trend radar requires regular scanning, defined update cycles, criteria for adding or removing entries, and communication routines. The challenge is to create a living knowledge resource, not a short-lived presentation.

Trust-and-Control Challenge - AI can speed up source discovery, summarisation, coding, clustering, monitoring and reporting. However, it may also produce hallucinated information, weakly supported synthesis, biased source selection, unclear classifications and overconfident language. Organisations must keep source traceability, distinguish evidence from interpretation, and retain human authority over trend approval. AI should support, not displace, foresight judgement.

These five challenges serve as a common framework throughout the working paper. The market scan identifies which tool types may address them; the test assesses selected configurations against them; and Chapter 5 translates the findings into recommendations.

2.3 Purpose and Scope of the Working Paper

This working paper examines how AI capabilities and AI-based tools can support the design, maintenance and communication of trend radars. It focuses on food processing SMEs while developing a framework that is relevant beyond this business sector. Its purpose is to identify AI-supported configurations that make trend radar work faster, more accessible and more continuous.

The paper addresses two **guiding questions**:

- (1) Which stages of the trend radar process can AI support, and through which tools and capabilities?
- (2) How can SMEs implement AI-supported trend radar processes efficiently?

The working paper thus treats the trend radar as a **socio-technical foresight system** rather than as a software product or visual output. This includes defining the purpose, framing the research, selecting sources, applying analytical methods, interpreting results, communicating findings, and maintaining the process over time.

Food processing serves as the application context because it combines consumer, technological, health, regulatory, sustainability and supply-chain developments. The case, *“Emerging Developments at the Interface of Nutrition, Beauty and Health with Relevance to Food Processing Companies”*, shows how the framework can be used to generate a first-draft trend radar and preliminary implications for food industry innovation.

The working paper includes a literature review, a market scan of AI-supported tools, a test of a trend radar case, and implementation recommendations. It does not assess the predictive accuracy of AI-generated trends or produce statistically generalisable rankings of AI tools. The paper treats transparency, source quality, bias, human review, and accountability as cross-cutting requirements.

2.4 Research Design and Paper Structure

This paper uses a practice-oriented design. It combines a literature review, a market scan of relevant tools, and a case study conducted through two AI-supported test configurations. Together, these elements link established foresight research with practical application.

The **literature review** establishes the conceptual basis. It briefly covers strategic foresight, horizon scanning, weak signals, emerging trends, and trend radars. It describes the conventional trend radar process, AI capabilities that may support it, and the benefits and limitations of computer-assisted foresight (Tranfield, Denyer and Smart, 2003).

The **market scan** identifies examples of generic AI assistants and research tools, agent-building and workflow-automation tools, enterprise AI platforms, and purpose-built foresight platforms. It applies the taxonomy defined in Section 4.1 and documents tool category, intended users, process coverage, AI capabilities, data and source model, source traceability, workflow automation, human-control options, collaboration and communication features and SME fit.

The **empirical component** consists of one common case, *“Emerging Developments at the Interface of Nutrition, Beauty and Health with Relevance to Food Processing Companies”*. The same research brief

is conducted through two contrasting approaches: Test 1 uses a generic, analyst-led, source-grounded workflow using *Perplexity*; Test 2 uses the purpose-built foresight and innovation platform *FIBRES*.

The test uses the following **assessment criteria** derived from the five challenges defined in Section 2.2:

- “Time to a first reviewable output” addresses the Time-to-Insight Challenge,
- “Source relevance, traceability, signal extraction, clustering and trend-description quality” address the Signal-to-Meaning Challenge,
- “Human effort and SME suitability” address the SME Capability Challenge,
- “Repeatability and support for updating” address the Living-Radar Challenge, and
- “Evidence inspection, workflow transparency and human approval” address the Trust-and-Control Challenge.

Across all elements, the working paper adopts a “**human-in-the-loop perspective**”: AI supports collection, processing, organisation, monitoring and communication, while humans remain accountable for framing, validation, interpretation, prioritisation and action.

Chapter 3 presents the conventional trend radar process. Chapter 4 defines the tool taxonomy, documents the market scan, analyses AI support across the workflow and presents the tests. Chapter 5 offers implementation recommendations, and Chapter 6 concludes.

3. Theoretical Framework

This chapter positions foresight within innovation management, explains horizon scanning and weak signals, and introduces trend radars as instruments for organising and interpreting future-oriented information. It then outlines data practices and the conventional radar process, which serves as the baseline for Chapter 4.

3.1 Foresight in Innovation Management

3.1.1 Innovation Management and Foresight

Innovation management concerns the systematic identification, development, selection and implementation of new or improved products, services, processes, technologies and business models (Jacobsen, 2025). It links an organisation's existing competencies and market position to potential future sources of customer value and competitive advantage. The innovation management framework (see Figure-1) used in this working paper positions foresight at the beginning of the innovation management process: it informs the identification of changing needs, technologies, market conditions and external constraints that may shape subsequent search, development and implementation activities.

Future-oriented intelligence does not replace customer research, technological expertise or strategic planning. It complements them by focusing attention on **developments that may not yet be visible** in market data or planning approaches. In this sense, foresight can support exploring future business fields, generating ideas, and critically reviewing established projects and assumptions (Rohrbeck and Gemünden, 2011).

The need for future-oriented intelligence arises because different approaches address different types of future questions. Table-1 distinguishes the approaches relevant to trend radar work. The order moves from relatively narrow approaches that estimate likely continuity to the broader, uncertainty-oriented practice of strategic **foresight**, of which **horizon scanning** is a core component.

Approach	Central Question	Output	Relevance to this Paper
Forecasting	What is most likely to happen?	Projection, estimate or probability statement	<i>Relevant as contrast:</i> does not adequately address discontinuities
Trend extrapolation	How might an observed pattern continue?	Time-based projection	<i>Relevant as contrast:</i> assumes that the underlying development remains stable
Foresight	Which plausible developments should inform current action?	Future intelligence, strategic options & shared understanding	<i>Central:</i> provides the overall process within which trend radars are used
Horizon scanning	What early signs of change require attention?	Source material, emerging issues and potential signals	<i>Central:</i> provides the information base for trend-radar work

Table-1: Future-oriented Approaches relevant to Trend Radar Work (Author's presentation)

The table does not include scenario development or roadmapping as they are not the primary subject of this working paper. **Scenario development** is used later to interpret the implications of trends and uncertainties together. **Roadmapping** may use foresight results to plan competencies development, but it is not a core stage of trend radar design or maintenance.

3.1.2 Foresight and Horizon Scanning

Foresight is a systematic, participatory process through which organisations gather, interpret, and use future-oriented intelligence to inform present-day decisions and collective action (Georghiou et al., 2008). In innovation management, its purpose is not to identify one predetermined future. It improves **preparedness** by making opportunities, threats, discontinuities, and capability requirements visible early enough to initiate action.

Uncertainty arises when the direction, timing, scale or consequences of future developments cannot be reliably derived from existing data. It becomes important where technological change, scientific developments, consumer values, environmental pressures, geopolitics, regulation, and business models interact. Under such conditions, approaches based mainly on historical trends or stable assumptions do not consider structural breaks or unexpected combinations of events.

Low-probability, high-impact developments are often described as **wild cards**; Taleb (2007) uses the related term “**black swans**” for highly consequential events that are difficult to foresee. Their existence does not make foresight useless. It underscores the need to consider alternative futures and improves the organisation’s capacity to detect and interpret early change.

In this process, **horizon scanning** provides a central function: it systematically searches for emerging developments and potential signals. Section 3.1.3 explains how weak signals are understood and why their interpretation is central to trend radar work.

3.1.3 Horizon Scanning and Weak Signals

Horizon scanning is the systematic search for emerging issues, developments, threats and opportunities that may have future relevance. It does not predict a single future. Instead, it collects and assesses current observations that may indicate possible change and provides evidence for subsequent interpretation and action (Rowe, Wright and Derbyshire, 2017).




Figure-2: Weak signals help to think about different Futures (Kuosa, 2010; Sitra, 2022)

Weak signals are relevant in this process. Ansoff (1975) used the term to describe early and incomplete indications of **potentially significant strategic developments**. Such signals are difficult to interpret because they are fragmented, uncertain and often only weakly connected to an organisation’s current

priorities. They may appear in scientific publications, patents, niche product launches, research projects, investment activity, policy consultations, new cultural practices or technological demonstrations.

Hiltunen (2008) distinguishes between an **observable issue** (= indication of change) and an **interpretation of its possible future significance**. This distinction is relevant because the observation itself may be well documented while its strategic meaning remains uncertain. A weak signal is therefore not simply a “small trend”. It is an observation that justifies interpretation and monitoring.

For this working paper, a **trend** is a more sustained and recognisable direction of change supported by multiple observations. The transition from observation to trend is not automatic: individual signals may fade, be reinterpreted or prove irrelevant, whereas a small number of signals may prove strategically important in a particular context.

Trend radars provide a way to organise this progression from observations to trends. Section 3.2 therefore examines their purpose, functions and architecture within foresight work.

3.2 Trend Radars in Foresight Work: Purpose, Functions and Architecture

A trend radar is a visual and analytical instrument for organising and communicating future-oriented information. Its **purpose** is to provide a shared, structured view of developments that justify attention in innovation management work. The trend radar is therefore not an end in itself; it turns observations into a basis for discussion, prioritisation and possible action.

Trend radars developed from technology-intelligence and corporate-foresight practice. Technology radars focus primarily on emerging technologies and their relevance for innovation strategy (Rohrbeck, Heuer and Arnold, 2006). In this working paper, the radar logic is extended to a broader set of change phenomena relevant to foresight, including social, economic, environmental, regulatory, consumer and business-model developments. This extension aligns with foresight approaches that consider multiple interacting drivers of change rather than technological factors in isolation (Saritas and Smith, 2011; Rohrbeck and Kum, 2018).

A trend radar fulfils five connected **functions**:

- Early warning by making possible change visible,
- Pattern recognition by bringing related observations together,
- Prioritisation through assessment dimensions,
- Tracking of developments over time, and
- Communication that supports shared understanding.

These functions are realised through the trend radar’s **architecture**: thematic sectors, a horizon or maturity logic, and a set of filters or visual cues. The intended decision context, rather than a standard graphic template, determines this architecture.




Figure-3: Exemplary Trend Radar Architecture (Jacobsen, 2025)

The radar provides the structure for the **analytical hierarchy** used in Section 3.3. Individual data points are collected through scanning; potentially relevant data points are interpreted as signals; sufficiently coherent signal clusters may be described as emerging trends; and trends are then considered for possible strategic implications. Thus, the hierarchy is both an analytical sequence and an information model for the radar:

Data points → signals → trends → strategic implications

This progression is not imperative. A data point may remain irrelevant, a signal may fade, and a cluster of signals may not justify being considered a trend. Computational methods can assist with retrieval, classification and semantic similarity, but human judgement is required to assess evidence and organisational relevance.

Trend Radar **sectors** may be organised by origin of change, for example political, economic, social, technological, environmental and legal factors (PESTEL); by strategic topic, such as personalised nutrition or packaging; or by value-chain impact, such as ingredients, manufacturing, distribution and consumption.

Trend Radar **horizons** may indicate time to impact, maturity or adoption stage, action priority or uncertainty. Further filters can capture impact, relevance, geography, affected business units or opportunity/threat orientation. For food processing SMEs, a manageable design may combine topic-based sectors with adoption-based horizons.

The trend radar is therefore both a communication tool and a structured representation of the underlying evidence and interpretation. It should provide an overview for decision-makers while allowing users to inspect detailed trend profiles, sources, uncertainty and update status. Section 3.3 explains how this information is collected, processed, and maintained.

3.3 Theory of the Trend Radar Development Process

3.3.1 Source Selection and Horizon Scanning

The trend radar development process begins by defining a **source strategy**. No single **source category** provides a complete view of future developments, so the choice should reflect the research brief and balance depth, breadth and reliability. Relevant **source categories** may include scientific and technical literature, patents, regulatory and policy material, industry reports and trade media, company and start-up activity, professional networks, selected consumer sources, and internal knowledge.

Different sources reveal different forms and stages of change. Research and patents may indicate technological possibilities before adoption; regulatory consultations may reveal future constraints; start-up activity may signal business-model experiments; and consumer sources may reveal changing language, practices or values. A documented source strategy makes both the evidence base and its blind spots visible.

Horizon scanning gathers potentially relevant **data points** through manual research, expert networks, databases, alerts, newsletters, web search, social listening or AI-supported retrieval. Each retained item should be recorded in a **source register** with its title, publisher or author, date, source type, link or identifier, a short evidence note, and an initial relevance assessment. This source register records provenance; it is distinct from the signal record created in the next stage.

3.3.2 Signal Extraction and Coding

Signal extraction converts potentially relevant data points into analytical entries. A **signal record** interprets an observation in relation to the research brief. It captures:

- What may be changing?
- Why it may matter?
- What evidence supports that interpretation?

It links to the source register but adds the radar sector, tags, geographical relevance and a statement of possible relevance.

This is not an entirely technical task. A product launch, publication or patent may be routine, a response to an established trend, a niche experiment or an indication of novel change. Analysts therefore combine **deductive coding**, using **predefined sectors** or value-chain categories, with **inductive tags** that capture unexpected connections. AI can assist through summarisation, preliminary classification, duplicate detection, semantic similarity analysis and tagging, but its outputs remain proposals for review (Mühlroth and Grottke, 2018).

3.3.3 Clustering and Trend Formation

Related signals may be formulated as **trends**. For example, observations concerning microbiome research, personalised dietary advice, at-home testing and beauty-from-within products may suggest a trend around data-enabled wellbeing nutrition. Clustering can be manual or computationally supported through semantic similarity, co-occurrence or network analysis.

A cluster of signals does not automatically constitute a trend. Analysts should consider whether supporting observations are sufficiently diverse and credible, whether they indicate a directional development rather than a temporary event, what connects them, how the cluster differs from existing entries, and which counter-signals or barriers require attention. The output is a small number of clearly defined, evidence-based trend candidates rather than a large collection of fashionable labels.

3.3.4 Evaluation, Prioritisation and Scenario Interpretation

After formulating “trend candidates”, organisations decide which require attention and how to position them on the trend radar. Management compares trend candidates against the organisation’s context, using a limited **set of criteria** such as potential impact, strategic relevance, time to impact or maturity, uncertainty, and opportunity or threat potential. Written justifications and source references are added to the ratings so scoring does not create a false impression of precision.

Scenario development follows this assessment, considering how trends and uncertainties could interact under different future conditions. It helps management to move from “**What is this trend?**” to “**What could this combination of developments mean for us?**” (Amer, Daim and Jetter, 2013; Rowe, Wright and Derbyshire, 2017).

3.4 Conventional Trend Radar Process and Data Practices

The preceding sections described the **theoretical components** of trend radar work. This section combines them into an end-to-end process. This process provides the baseline against which Chapter 4 assesses AI-supported approaches.




Figure-4: Conventional Trend Radar Process (Author’s own visualisation)

3.4.1 Define Search Areas

The first stage is to define why the radar is being developed and which decisions it is intended to support. Possible purposes include strategic planning, innovation tracking, technology intelligence, operational readiness, risk monitoring, opportunity identification, and communication of future developments.

The subsequent **research brief** specifies:

- The radar title and central research question,
- Intended users and decision-makers,
- Strategic purpose and expected use,
- Geographic, sectoral, and value-chain scope,
- Time horizon,
- Radar sectors and horizon logic,
- Relevant source categories,
- Inclusion and exclusion criteria,
- Required output formats,
- Update and review interval,
- Roles and approval responsibilities.

This stage determines **what is considered relevant**. A broadly defined scope produces a large amount of interesting but unusable information. A narrowly defined scope overlooks important adjacent developments. The research brief should provide sufficient focus while allowing unexpected signals to be identified and assessed.

3.4.2 Define Source Strategy

The second stage defines the trend radar's underlying **information structure**, which should also be reflected in the research plan. This includes sectors, horizon categories, assessment dimensions, tags, filters, and visual logic. It also defines the source strategy.

A **source strategy** balances depth and breadth by identifying priority source categories, search channels, keywords, source-specific collection intervals, and quality criteria. The strategy distinguishes between **recurring source categories**, such as scientific alerts, patent feeds, regulatory updates, industry newsletters, and selected media outlets, and **ad hoc source categories**, such as event reports, expert recommendations, or newly discovered databases.

The source strategy should also address potential **bias**. A radar based only on English-language digital sources may underrepresent local, non-English, non-digital, or less visible developments. A radar based only on trade media may push commercially promoted topics. A radar based only on scientific literature may overlook changing consumer practices or emerging business models. **Source diversity** and periodic source review are important.

3.4.3 Identify Signals

The third stage consists of **horizon scanning**, more specifically **collecting relevant data points**. Data is collected manually, semi-automatically, or automatically, depending on available tools. Regardless of the collection method, retained items are recorded in a structured **source register**.

The source register includes, at a minimum:

- Source title,
- Author, publisher, or organisation,
- Publication date and retrieval date,
- Source type,
- Link or document location,
- Short summary or quotation,
- Relevant keywords and tags,
- Initial relevance assessment,
- Responsible analyst.

At this stage, the objective is not to determine whether each item represents a trend. The objective is to capture potentially **relevant observations** without discarding “unusual” material. However, analysts exclude duplicates, unreliable sources, and material outside the agreed scope.

3.4.4 Categorise Signals

The fourth stage transforms collected information (= data points from observations) into a structured **signal base**. Analysts review data points, assess their relevance, add metadata, assign sectors and tags, and identify potential relationships between items. Where appropriate, the analysis may distinguish between signals, signal clusters, trends, and megatrends.

The assessment is documented so users understand why an item was included or excluded. For each signal, the process should record the evidence, interpretation, assigned classification, level of uncertainty, and any open questions. This **distinction between evidence and interpretation** matters, as it lets reviewers revisit earlier assumptions when new information becomes available.

Related **signals** are clustered. At this point, organisations must resist labelling every cluster as a trend. Some clusters will remain weak or fragmented. Others may merge with existing themes or disappear over time.

3.4.5 Build Trend Radar

The fifth stage involves deciding which groups of signals to classify as **trends** and how to position them on the radar. Trend descriptions should avoid vague or promotional language. They should explain what is changing, which drivers support the interpretation, and why the development matters to the organisation.

A useful **trend profile** includes:

- Trend title and short definition,
- Related signals and source links,
- Description of the underlying development,
- Key drivers,
- Relevant stakeholders or affected value-chain stages,
- Time horizon or maturity level,
- Potential impact and uncertainty,
- Opportunity and threat considerations,
- Date of the latest review and responsible owner.

Radar positioning translates these assessments into a **visual logic**. For example, a trend may be positioned by sector and maturity level, while marker size indicates estimated impact. The visual position should always be interpreted as judgement, not objective fact.

3.4.6 Create Narratives

The sixth stage connects trend information to strategic meaning. A trend radar is useful when it helps management ask better questions and consider possible responses. This may include developing trend narratives, constructing future tables, creating scenarios, or identifying strategic options.

Trend narratives explain why a development matters and how it may evolve. They can be tailored to different audiences. A management narrative may focus on strategic implications, risks, opportunities, and decisions. A technical narrative may focus on scientific feasibility, process requirements, intellectual property, or regulatory status. A consumer-oriented narrative may focus on behaviour, values, needs, and market segmentation.

Scenarios then combine multiple trends and uncertainties to create plausible future contexts.

3.4.7 Disseminate Results

The final stage concerns dissemination, monitoring, and continuous improvement. A trend radar should not be treated as a one-off report. It should be a **living resource** that is regularly reviewed, updated, and used.

A practical maintenance model includes:

- Continuous or weekly scanning of selected sources,
- Monthly review of newly collected signals,
- Quarterly updates of priority trends and radar positions,
- Annual review of the radar's purpose, sectors, source base, and assessment criteria,
- Ad hoc alerts for high-priority developments, rapid changes, or critical regulatory signals,
- Feedback mechanisms for users and experts.

This trend radar process demonstrates **why AI may be valuable**: several stages are labour-intensive, repetitive, data-rich, and suited to computational support. At the same time, the process also shows why **AI cannot replace foresight practice**. Framing the strategic question, selecting credible sources, interpreting weak signals, evaluating relevance, handling uncertainty, and translating insights into action require human involvement.

The following chapter examines how different AI tools and configurations can support these stages while maintaining transparency, governance, and meaningful organisational use.

4. AI-supported Trend Radar Approaches

4.1 Classification of relevant AI Tools

In this working paper, AI refers to computational capabilities that assist the collection, processing, interpretation, monitoring and communication of foresight information. Relevant capabilities include:


- Generative AI,
- Retrieval and source-grounded synthesis,
- Natural-language processing (NLP),
- Machine-learning methods (ML), and
- Agentic workflows that coordinate multi-step tasks.

These capabilities overlap. Generative AI can draft research plans, summaries and trend descriptions. Retrieval-augmented systems combine language models with external documents or sources, enabling evidence inspection. NLP and ML can support classification, semantic similarity, clustering and prioritisation. Agentic workflows can connect retrieval, processing, review and reporting across a defined sequence of tasks.

For the identification of potentially useful tools in the market scan, AI-supported tools are classified along **two dimensions**:

The **first dimension** distinguishes generic AI assistants and research tools, generic agent-building and workflow-automation tools, enterprise AI-development and governance platforms, and purpose-built foresight and trend-intelligence platforms.

The **second dimension** distinguishes relatively simple, primarily single-user or single-task workflows from relatively complex workflows that are multi-step, automated, integrated or collaborative. Complexity in this context describes a typical pattern of use rather than a hierarchy of quality.



Note: Placement by primary workflow orientation. Tools may support more than one quadrant depending on configuration, integrations, user skills and organizational use.

Figure-5: AI tool taxonomy (Author's own visualisation)

Section 4.2 applies this taxonomy to selected tools. It is developed for this working paper and combines the conventional trend radar process described in Chapter 3 with the five practical challenges identified in Section 2.2 and concerns in AI-supported foresight, including source provenance, transparency, human

oversight, workflow integration and organisational fit (Mühlroth and Grottke, 2018; Friedewald et al., 2025).

The purpose is not to identify the tool with the most features, but to assess which configurations can **address the challenges** of time to insight, signal-to-meaning, SME capability, continuous maintenance, and trust and control.

4.2 Market Scan of AI-Supported Trend Radar Tools

The desk-based market scan applies the taxonomy in Section 4.1 to illustrative tools. It uses vendor information, documentation, and independent sources to identify tools using **dimensions** such as tool category, intended users, process coverage, AI capabilities, data and source model, source traceability, workflow automation, human control options, collaboration and communication, and SME fit (see Annex details). The scan does not provide an independent evaluation or ranking of the tools. Its purpose is to develop the subsequent test case and justify selecting the two configurations examined in Section 4.3.

Included tools have publicly documented functionality relevant to at least one stage (see chapters 3.4.1 to 3.4.7) of the trend radar process and represent a distinct workflow pattern:

Generic AI Assistants - Generic assistants have low entry barriers and support research briefs, keyword generation, source exploration, summarisation, preliminary coding and draft trend profiles. *ChatGPT*, *Claude* and *Gemini* are general-purpose generative assistants; *Perplexity* is more source-oriented and supports source discovery and source-linked summaries. These tools do not themselves provide a robust radar environment, so users must establish the source register, classification logic, review process and update routine.

Generic Agent-Building and Workflow-Automation Tools - *AgentGPT*, *Lindy*, *CrewAI* and *AutoGen* represent repeatable, partially automated workflows. They can connect source monitoring, initial filtering, classification, routing and reporting. Their relevance lies in recurring operational routines rather than autonomous strategic judgement. They are useful once source criteria, taxonomies and review roles are established.

Enterprise AI-Development and Governance Platforms - Enterprise platforms such as *IBM watsonx* support bespoke applications that integrate internal data, retrieval-based systems, machine-learning models, and governance mechanisms. Their advantages in integration, security, and scale must be considered against substantial technical and organisational requirements; SMEs are more likely to access them through associations, research institutions, or shared-service models.

Purpose-Built Foresight and Trend-Intelligence Platforms - Purpose-built platforms provide foresight-specific information models, workflows, visualisation and collaboration. *Nextatlas* focuses on consumer intelligence; *Futures Platform* combines curated trends, signals and scenarios; *FifthRow* offers AI-native research and strategy workflows; and *Itonics* combines trend and technology intelligence with collaborative evaluation, portfolio and roadmapping functions. *FIBRES* is relevant because it supports organising trend data, radar visualisation, team communication, and ongoing maintenance of a shared trend repository. Such platforms support structured work but do not replace source validation or strategic interpretation.

For the test (section 4.3), the market scan identified two contrasting configurations: *Perplexity* represents a generic, analyst-led, source-grounded workflow; *FIBRES* represents a purpose-built trend-radar platform with information-management, visualisation, and maintenance functions. They are not selected to establish a universal ranking, but to compare two plausible SME pathways for applying the same research brief.

4.3 Test Case: “I Feel Good”

The case, “Emerging Developments at the Interface of Nutrition, Beauty and Health with Relevance to Food Processing Companies”, was examined through two exploratory AI-supported configurations.

Test Configurations & Procedure - Test 1 uses *Perplexity* in an analyst-led, source-oriented workflow; Test 2 uses *FIBRES* to configure and automatically populate a purpose-built trend radar. Both tests begin from the common research brief in Table-3, but they do not follow identical operational procedures or reach the same review status. The comparison therefore examines what each documented configuration produces and what human work remains necessary.

The intended analytical sequence is to define the brief, identify and review sources, extract signals, examine clusters, formulate trends, position selected entries, and prepare management communication. In Test 1, the analyst explicitly guides this sequence and records human decisions at successive approval gates. In Test 2, the brief was entered into *FIBRES* and its automation generates research topics, retrieves material, synthesises phenomena, and populates a radar in one automated run. The supplied record ends with an instruction to review the AI-generated content; it does not document completed source, trend, or radar-position approvals. Consequently, the *FIBRES* artefacts are assessed as an automatically populated first draft, not as an approved final radar.

Research Questions - The case operationalises the two guiding questions in Section 2.3:

- (1) Which stages of the trend radar process can AI support, and through which tools and capabilities?
- (2) How can SMEs implement AI-supported trend radar processes efficiently?

Comparative Assessment Criteria - The comparative assessment framework is researcher-constructed. It translates the conventional trend-radar process described in Section 3.4, the five challenges defined in Section 2.2, and the two guiding questions in Section 2.3 into observable, documentable criteria.

The framework draws on research on computer-supported weak-signal and trend detection, which highlights the relevance of data collection, text analysis, classification, and clustering while also emphasising the interpretive and uncertain character of foresight outputs (Mühlroth and Grottke, 2018). It further reflects the “human-in-the-loop” perspective: automated support may increase analytical capacity, but source evaluation, contextual interpretation and approval remain human responsibilities (Friedewald et al., 2025).

Guiding Question	Challenge addressed	Criterion	Practical assessment question
1: Which stages of the trend radar process can AI support, and through which tools and capabilities?	Time-to-Insight	Time to first usable output	How long does it take from research brief to a source-linked, first-draft trend radar?
	Time-to-Insight	Research-plan quality	Does the tool produce a coherent and usable research plan, w/ keywords, source categories, sectors and analytical steps?
	Signal-to-Meaning	Source relevance	Are the sources relevant to the defined topic, geography, time horizon and innovation focus?
	Signal-to-Meaning; Trust-and-Control	Source traceability	Can a reviewer inspect the original source, date, source type and the connection between evidence & the proposed trend?
	Signal-to-Meaning	Signal usefulness	Do the signals describe potential changes rather than duplicates, generic claims, or irrelevant material?
	Signal-to-Meaning	Trend-description quality	Are trend titles, descriptions, evidence, uncertainty and implications clear, specific and supported by linked evidence?
	Signal-to-Meaning	Radar readiness	Can the output be converted into a visualisable radar with sector, horizon/ maturity, profile, evidence links & filters?
	SME Capability	Human effort	How much analyst and reviewer work is required for prompting, source checking, cleaning, deduplication, classification, correction and final approval?
2: How can SMEs implement AI-supported trend-radar processes efficiently?	Trust-and-Control	Transparency of evidence and workflow	Can a reviewer reconstruct how a trend entry was produced, including prompts or workflow, sources, and human changes?
	SME Capability	Skills and configuration effort	What skills, training, set-up work, workflow design and tech support are required before a team can use the approach?
	SME Capability	Cost and access conditions	What subscriptions, licences, paid source access, platform accounts or technical infrastructure are required?
	Living-Radar	Repeatability	Can the workflow be repeated with a similar process, documentation standard and human effort when new information is available?
	Living-Radar	Support for ongoing updating	Can the approach support recurring source collection, re-evaluation of signals, updates to trends and versioned radar maintenance?
	SME Capability; Trust-and-Control	SME suitability	Considering effort, skills, cost, transparency, maintenance requirements and output quality, is the approach proportionate for a food-processing SME?
	Trust-and-Control	Human authority	At which points can human users inspect, challenge, revise, approve or reject AI-supported outputs?

Table-2: Assessment Criteria (Author's own presentation)

Common Case and Research Brief - Both tests use the same case: “*Emerging Developments at the Interface of Nutrition, Beauty and Health with Relevance to Food Processing Companies*”. The configurations are initiated with the same case and research brief in Table-3. This provides a shared reference for assessing relevance, required outputs, and human responsibilities. It does not make the tests operationally equivalent: the *Perplexity* workflow involves repeated analyst-directed source expansion and staged approval, whereas the *FIBRES* run uses platform-generated topics and automated radar population without a recorded post-generation approval cycle. Differences in the resulting artefacts must therefore be interpreted together with differences in configuration and review effort.

Element	Common specification for both tests
Case / radar title	Emerging Developments at the Interface of Nutrition, Beauty and Health with Relevance to Food Processing Companies
Strategic purpose	Identify, analyse and organise weak signals and trends relevant to food product innovation at the intersection of nutrition, beauty and health
Intended use	Provide a source-linked, transparent first-draft visual trend radar to support ideation of future product opportunities
Primary users	Innovation managers and senior leaders in German food processing SMEs
Secondary users	Product-development, marketing, sustainability, regulatory, strategy and research personnel in German food processing SMEs
Geographic scope	North America and Europe, supplemented by selected global developments where strategically relevant
Sectoral scope	The interface of (1) food & health, (2) food & beauty, (3) beauty & health, and (4) food & beauty & health, with relevance to food product innovation
Value-chain scope	Food & beverage product development, food & beverage product manufacturing, and food & beverage consumption.
Time horizon	Approximately ten years from now on
Horizon logic	Trends are assessed by maturity/adoption stage: emerging, growing, maturing and mainstream. Where possible, an indicative time-to-impact assessment is recorded
Core radar sectors	Consumer & Market Practices (Social); Technology & Data (Technological); Sustainability (Environmental); Regulation (Legal)
Analytical hierarchy	Data points → signals → trends → strategic implications
Relevant source categories	1st priority: Consumer & cultural developments; start-up & investment activity; company & product announcements; industry reports & trade publications 2nd priority: Scientific & technical literature; patents; regulatory & policy material
Source-quality requirements	Prioritise identifiable & inspectable sources. Record title, author/publisher, date, source type and link or identifier. Distinguish between primary sources, secondary sources and promotional material
Inclusion criteria	Include material that is: relevant to the defined topic, geography and time horizon; provides identifiable evidence of a possible change; and has potential implications for food product innovation
Exclusion criteria	Exclude: duplicate entries; material without an identifiable source; content outside the defined scope; unsupported AI-generated claims; and observations with no plausible connection to food product innovation
Required output: source register	A documented register of relevant retained sources, including source metadata, short evidence notes and initial relevance assessment
Required output: signal records	A structured set of potential signals linked to original sources and assigned to a preliminary sector or category
Required output: signal clusters	Proposed groups of related signals, including a brief explanation of the suggested relationship and any significant uncertainty or counter-signals
Required output: trend profiles	For each selected trend: title, concise description, supporting signals and sources, relevant sector, maturity or horizon assessment, possible implications and uncertainty statement
Required output: first-draft radar	A visual representation of selected trends organised by core sectors and horizon logic
Required output: management summary	A short, accessible summary of the most relevant trends, uncertainties and preliminary implications for German food processing SMEs
Required output: methodological record	Principal prompts, search logic, workflow settings, manual interventions, corrections and unresolved issues
Review and update rhythm	For the initial test, produce one first-draft radar. For continuing use, assume a monthly review of new signals, quarterly update of priority trends and annual review of scope, sectors and sources
Required roles and approval responsibilities	Human reviewers approve the research brief, source inclusion, signal-to-cluster relationships, trend formulation, radar positioning and final communication. AI tools provide proposals and analytical support only

Element	Common specification for both tests
Test 1 configuration	<i>Perplexity</i> supports source-oriented research and source-linked preliminary summaries; supports research planning, synthesis, coding structures and draft trend profiles
Test 2 configuration	<i>FIBRES</i> . The platform supports structuring, linking, reviewing, visualising, communicating and maintaining trend-radar content
Assessment approach	Qualitative and descriptive comparison using the criteria specified in table-2, Section 4.3
Assessment dimensions	Time to first reviewable output; research-plan quality; source relevance and traceability; signal usefulness; clustering quality; trend-description quality; radar readiness; human effort; repeatability; support for updating; and SME suitability
Time stamp and tool conditions	Record the date of each test, tool version or access mode where identifiable, subscription level, source-access conditions and material changes to the workflow

Table-3: Research brief (Source: Author's own presentation)

4.3.1 Test 1: Analyst-led Workflow using *Perplexity*

The *Perplexity* test began with a source-inclusion proposal containing 17 candidate sources. The analyst judged this coverage insufficient and supplied several targeted lists of trade, start-up, industry, consumer, and scientific sources. Successive searches expanded the documented register to 132 source records. This expansion is a substantive part of the result: **the breadth of the evidence base cannot be attributed to the initial open-ended search.**

The workflow proceeded through human-review gates. Source decisions preceded signal extraction; signal and cluster decisions preceded trend formulation; and trend decisions preceded radar positioning and management communication. Following review, the record contained 38 approved signals, 11 clusters, and ten positioned trends. The *Perplexity* output distinguished time-to-impact, maturity, and assessment confidence and disclosed material evidence gaps.

The analysis distinguished commercial activity from demonstrated effects. It retained weak signals with markers, did not claim verified environmental performance, identified unreliable regulatory evidence and a German consumer-evidence gap, and did not formulate a trend asserting a demonstrated ingestible-to-appearance effect. An English management summary was produced. At the end of the session, approval of its content and release remained pending. The *Perplexity* result is thus an **approved first-draft radar accompanied by a proposed management communication.**




Figure-6: Perplexity Output

4.3.2 Test 2: Automated Population using FIBRES

For the *FIBRES* test, the research brief was entered as a radar definition with four sectors: Consumer & Market Practices, Technology & Data, Sustainability, and Regulation. Adoption horizons were configured from Mainstream at the centre to Emerging at the outer rim. *FIBRES* expanded the purpose statement and generated 153 research keywords across shared and sector-specific topics. The keyword set included directly relevant concepts such as beauty-from-within, collagen, microbiome, and claims substantiation, as well as broader terms on packaging, biodiversity, traceability, and general wellness.

The platform reported 3,200 search results, 3,018 collected articles, 959 potential phenomena, 58 synthesised phenomena, and 77 links between phenomena. Its generated report separately described 857 unique articles. Because the supplied record does not reconcile the 3,018 collected articles with the 857 unique articles, these figures are distinct system-generated counts and should not be interpreted as measures of the final evidence base. *FIBRES* produced a populated radar, network visualisation, report,

and exportable records. This demonstrates rapid structuring and visualisation, although elapsed processing time was not recorded in a form suitable for direct comparison with Test 1.




Figure-7: Automatically populated FIBRES Draft Radar, prior to documented Human Review

The automated first population highlights a relevance-control problem. The material includes entries on food-relevant topics alongside examples from construction adhesives, apparel, cosmetics packaging, and other domains whose connection to the food-product decisions is unclear. Its source material combines research and specialist trade coverage with company announcements, press-release services, and market forecasts. Links and descriptions enable subsequent inspection, but the supplied workflow does not document source-by-source inclusion decisions or a human-approved separation of observations, signals, clusters, and trends before radar positioning.

The system itself stated that its agents had performed the work and requested human review. The resulting visualisation should therefore be called an **automatically populated draft radar**; it is not an analytical equivalent to the reviewed *Perplexity* radar.

4.3.3 Cross-case Assessment

The two configurations support different parts of the trend radar process particularly well. *Perplexity* supported a source-linked, analyst-directed progression through defined approval gates, but the analyst had to enlarge the source frame and maintain the workflow and artefacts. *FIBRES* provided integrated radar configuration, automated scanning and synthesis, exportable item records, and immediate network and radar views; its first population required further relevance filtering and expert review before management use. In neither test does the number of retrieved articles or plotted entries establish the quality of the resulting trends.

Against the working paper's five challenges, the contrast is clearest for time-to-insight, signal-to-meaning, and trust-and-control. *FIBRES* reduced the apparent steps between configuration and visual output, whereas *Perplexity* exposed the decisions between evidence and an approved radar. *FIBRES*' broad discovery expanded coverage but also retrieved material outside the defined scope; *Perplexity*'s more selective record depended on extensive human source prompting. A persistent platform offers a plausible basis for ongoing maintenance, while the *Perplexity* workflow supplies a clearer documented approval trail. The supplied records do not support either a comparative person-hour estimate or actual performance across repeated radar-update cycles.

These findings support a **limited conclusion**: AI can assist source discovery, synthesis, classification, relationship mapping, and visual preparation, but a first-draft radar becomes decision-ready only when relevance, evidence quality, interpretation, and positioning have been reviewed against its intended use. Chapter 5 translates this conclusion into recommendations for SMEs.

5. Recommendations for Implementation

The preceding chapters describe trend radar work as a sequence of decisions about purpose, sources, interpretation, positioning, and use. The two configurations illustrate different ways AI can support that sequence. In the *Perplexity* test, an analyst directs a source-oriented workflow through documented review gates. In the *FIBRES* test, a configured platform rapidly collects material and populates a radar. **Neither test establishes a generally superior tool.** Together, they suggest that an SME should implement an AI-supported trend radar as a supervised research and learning process, rather than accepting an automated graphic as its finished foresight output. The recommendations below follow the five practical challenges introduced in Section 2.2.

5.1 Define Purpose, Decision Context, and Governance

Implementation should begin with the decisions the radar is intended to inform. A food-processing SME might use it to identify product-development questions, monitor ingredient technologies, examine changing consumer needs, or identify issues that require regulatory review. These purposes are related, but they do not require identical evidence, horizons, or outputs. The research brief should therefore name its users, strategic question, geographic and value-chain boundaries, time horizon, radar sectors, inclusion rules, and expected decisions before a search or automated scan begins.

The tests show why this **specification must remain under human control**. The common brief defined relevance as developments at the nutrition-beauty-health interface with implications for food-product innovation. In *FIBRES*, automated topic generation added a wide vocabulary covering, among other subjects, packaging, traceability, biodiversity, cosmetics, and general wellness. Some adjacent subjects may prove useful, but a keyword match alone does not establish relevance to the SME's decision. Generated search topics and radar categories should therefore be reviewed as part of the research plan, rather than accepted as a neutral translation of the brief.

Management should specify at least an analyst, subject-matter reviewers, and the person authorised to approve management communication. The *Perplexity* test stopped for decisions on source inclusion, signals and clusters, trends, radar positions, and final communication. The documented *FIBRES* run requested human review after automatically plotting its findings; the supplied record does not document equivalent approvals. An SME need not reproduce five meetings, but it should **record who approved the evidence base, the selected trends, and their communication**. A platform instruction to review all content does not, by itself, establish a review procedure.

5.2 Build the Data and Source Foundation

A reliable trend radar requires a source strategy in addition to search terms. The *Perplexity* test began with 17 source candidates; successive analyst-supplied lists substantially widened the material considered and produced a register of 132 approved records. This does not mean that a larger register is necessarily better. It shows that **open-ended discovery alone missed specialist publications** and that source selection affected which developments and counter-developments became visible. SMEs should **define a permanent source collection** including trade and company information, consumer research, scientific evidence, and primary regulatory material.

Automated breadth presents the opposite risk. *FIBRES* reported a large automated retrieval and a trend radar containing 58 synthesised phenomena, but the supplied export also includes items concerning construction adhesives, apparel traceability, and other topics without a clear connection to food-product innovation. These may be relevant observations in their own thematic areas; their presence does not

make them suitable evidence for this trend radar. The first filtering question should be: “*What decision about food products could this item inform?*” Material that cannot answer it should be excluded.

Every retained source should include its title, publisher, source type, publication and retrieval dates, original location, evidence note, relevance rationale, and review status. Analysts should distinguish a company’s announcement from independently observed adoption, a forecast from measured demand, and a laboratory result from demonstrated benefit in a finished food product. Publication dates and source claims also require checking. This register is not for administrative completeness: it lets a reviewer reconstruct why a signal or trend was accepted and revise it as better evidence emerges.

5.3 Combine AI Automation with Expert Review

AI is well suited to generating search suggestions, summarising documents, proposing tags, detecting possible duplicates, identifying related records, drafting trend profiles, and preparing visual or written first drafts. The two tests illustrate different strengths. *Perplexity* supported an incremental, analyst-directed chain from sources through signals and clusters to an approved first-draft radar. *FIBRES* generated a configured radar, phenomenon analyses, links, a network view, and a report with considerably less documented step-by-step intervention. These capabilities are useful, but volume and speed alone do not ensure analytical quality.

A practical workflow should separate at least five analytical objects: **source observations, interpreted signals, signal clusters, trend candidates, and strategic implications**. AI may propose links between them, but each transition needs a review rule. A source does not automatically establish a signal; several related signals do not automatically establish a trend; and a trend does not automatically justify an investment. The *Perplexity* review illustrates the importance of these distinctions: one proposed signal was rejected, three weak signals were retained with markers, and one candidate was reclassified from a positioned trend to an illustrative case.

Review should be linked to the claim. A product launch may be sufficient evidence of commercial activity, but an efficiency claim requires different evidence. A source on life-cycle-assessment methods does not establish the environmental performance of a particular ingredient. The analyst should record counter-evidence and missing primary material alongside supportive findings. In Case 1, the absence of a proofed ingestible-to-appearance effect, the lack of evidence to substantiate an environmental-performance claim for the assessed case, and the German consumer-data gap remained visible in the final draft.

For *FIBRES*, the initial automated population should consequently be treated as a **candidate collection**. Before management use, reviewers should remove off-scope items, inspect source claims, clarify whether an entry is a signal, technology, cluster, or trend, and approve its sector and adoption position. This recommendation concerns the documented case; it does not claim the platform lacks functions to carry out this review.

5.4 Maintain a Living Trend Radar

A first trend radar establishes a baseline; it does not establish a completed foresight capability. Maintenance requires an owner and a process for collecting new sources, reconsidering existing signals, and revising or archiving trends. The research brief proposes monthly reviews of new signals, quarterly reviews of priority trends, and annual reviews of the scope and source base. An SME can adapt those intervals to its resources, but should also define ad hoc triggers, such as an important scientific review, a material market development, or a relevant regulatory decision.

The maintenance record should show what changed and why: new or removed sources, merged signals, changes to maturity or time-to-impact judgements, and the reviewer responsible. A platform and network

view, such as *FIBRES*, makes relationships and revisions easier to manage than a sequence of separate documents. Conversely, constancy has limited value if unreviewed material builds up. The organisation should therefore judge maintenance support by whether it can preserve a trustworthy history of decisions, not simply by whether it can store or display more items. The present tests demonstrate initial production but do not yet measure performance over repeated update cycles.

5.5 Design for Visualisation, Communication, and Decision Integration

The trend radar should provide a readable overview while allowing users to inspect the source behind each entry. Sectors, horizons, colours, and other visual variables require stable definitions. In the *Perplexity* case, the reviewed radar used rings for indicative time-to-impact, colour for maturity, and marker opacity for assessment confidence. *FIBRES* used adoption stages as its radar horizons. Both are valid design choices, but they encode different analytical questions.

A visually populated sector is not necessarily analytically well supported. Case 1 explicitly disclosed limited regulatory evidence and did not use its retained source register to make an environmental-performance claim for the assessed food-product case. The *FIBRES* first population displayed substantially more sustainability-related material, some of it outside the food-product scope. The appropriate response is neither to leave every sector empty nor to populate every sector automatically. It is to show only reviewed entries and to state where coverage, relevance, or confidence remains limited.

Finally, dissemination should connect each selected trend to a management question: whether to monitor, commission evidence, test a formulation, explore a partnership, or reconsider a product assumption. A short management narrative should distinguish observed developments from possible implications and retain material uncertainties in the main text.

For an SME, **the most useful trend radar is not the one with the greatest number of signals. It is the one that helps a small team revisit evidence and make better-informed innovation decisions without hiding its limits.**

6. Conclusion

This working paper examined how AI can support the design and maintenance of trend radars for food-processing SMEs. It combined a conceptual framework, a market scan, and two exploratory tests using the same case on nutrition, beauty, and health. The aim was not to rank the tools, but to identify useful AI capabilities and the conditions under which their outputs can inform innovation decisions.

The first guiding question asked which stages AI can support. The tests illustrate its use in research planning, source discovery, summarisation, signal extraction, clustering, trend description, relationship mapping, visualisation, and communication. *Perplexity* supported an analyst-directed, source-linked workflow with explicit review gates. *FIBRES* rapidly generated a populated radar, network view, and report from a configured brief. **Neither configuration made strategic interpretation or evidence appraisal unnecessary.**

The second question asked how SMEs can implement these capabilities efficiently. The tests suggest that a decision-oriented brief, controlled source and topic frames, inspectable evidence, and defined approvals matter more than the volume of material retrieved. *Perplexity*'s initial source coverage required substantial human-led expansion, but its review process made trend-status changes visible. *FIBRES* shortened the path to a visual first draft, but included material whose relevance to the food-product case had not been established. **A populated trend radar is therefore not necessarily decision-ready.**

These findings are limited to one illustrative case and two differently reviewed workflows. The tests did not establish comparable person-hours or costs, predictive accuracy, performance across update cycles, or effects on management decisions. The *Perplexity* trend radar was approved as a first draft, while its management summary awaited final content approval; the documented *FIBRES* radar remained an automated draft awaiting human review. **The outputs should not be compared as equally validated products.**

Future tests should document review effort and corrections over repeated updates, including a human-reviewed *FIBRES* iteration. For SMEs, the principal opportunity is not an automatically completed picture of the future, but a manageable, living evidence base that helps people recognise change, examine uncertainty, and make better-informed innovation choices.

References

- Amer, M., Daim, T.U. and Jetter, A. (2013) 'A review of scenario planning', *Futures*, 46, pp. 23-40.
- Asioli, D., Aschemann-Witzel, J., Caputo, V., Vecchio, R., Annunziata, A., Næs, T. and Varela, P. (2017) 'Making sense of the "clean label" trends: A review of consumer food choice behaviour and discussion of industry implications', *Food Research International*, 99, pp. 58-71.
- Ansoff, H.I. (1975) 'Managing strategic surprise by response to weak signals', *California Management Review*, 18(2), pp. 21-33.
- Bigliardi, B. and Galati, F. (2013) 'Innovation trends in the food industry: The case of functional foods', *Trends in Food Science & Technology*, 31(2), pp. 118-129.
- Friedewald, M. et al. (2025) 'New perspectives for data-supported foresight: The hybrid AI-expert approach', *Futures & Foresight Science*, 4, e99.
- Georghiou, L., Cassingena Harper, J., Keenan, M., Miles, I. and Popper, R. (eds.) (2008) *The handbook of technology foresight: Concepts and practice*. Cheltenham: Edward Elgar.
- Gershman, M., Bredikhin, S. and Vishnevskiy, K. (2016) 'The state of foresight in small and medium enterprises: Literature review and research agenda', *Foresight*, 18(4), pp. 404-417.
- Hiltunen, E. (2008) 'The future sign and its three dimensions', *Futures*, 40(3), pp. 247-260.
- Mühlroth, C. and Grottke, M. (2018) 'A systematic literature review of mining weak signals and trends for corporate foresight', *Journal of Business Economics*, 88(5), pp. 643-687.
- Jacobsen, B.P., 2025. *Introduction to Innovation Management: Basics, Framework, Creativity*. Gross Grönau: NMC Edu.
- Kuosa, T., 2010. Futures signals sense-making framework (FSSF): A user-friendly tool for identifying and categorizing signals. *Futures*, 42(1), pp.42-50.
- Rohrbeck, R. (2010) *Corporate foresight: Towards a maturity model for the future orientation of a firm*. Heidelberg: Physica-Verlag.
- Rohrbeck, R. and Gemünden, H.G. (2011) 'Corporate foresight: Its three roles in enhancing the innovation capacity of a firm', *Technological Forecasting and Social Change*, 78(2), pp. 231-243.
- Rohrbeck, R., Heuer, J. and Arnold, H.M. (2006) 'The technology radar: An instrument of technology intelligence and innovation strategy', in *Proceedings of the 3rd IEEE International Conference on Management of Innovation and Technology*. Singapore, 21–23 June. Piscataway, NJ: IEEE, pp. 978-983.
- Rohrbeck, R. and Kum, M.E. (2018) 'Corporate foresight and its impact on firm performance: A longitudinal analysis', *Technological Forecasting and Social Change*, 129, pp. 105-116.
- Rowe, E., Wright, G. and Derbyshire, J. (2017) 'Enhancing horizon scanning by utilizing pre-developed scenarios: Analysis of current practice and specification of a process improvement to aid the identification of important "weak signals"', *Technological Forecasting and Social Change*, 124, pp. 111-124.

Saritas, O. and Smith, J.E. (2011) 'The big picture—trends, drivers, wild cards, discontinuities and weak signals', *Futures*, 43(3), pp. 292-312.

Sitra (2022), *Weak Signals 2022: Stories about the future*. Sitra Studies 203. Helsinki: Sitra.

Taleb, N.N. (2007) *The black swan: The impact of the highly improbable*. New York: Random House.

Vecchiato, R. (2012) 'Environmental uncertainty, foresight and strategic decision making: An integrated study', *Technological Forecasting and Social Change*, 79(3), pp. 436-447.

Vecchiato, R. (2023) 'Foresight for small and medium enterprises: Overcoming resource and capability constraints', *IEEE Engineering Management Review*, 51(3), p. 19-25.

Tranfield, D., Denyer, D. and Smart, P. (2003) 'Towards a methodology for developing evidence-informed management knowledge by means of systematic review', *British Journal of Management*, 14(3), pp. 207-222.

Annex

Tool / framework	Tool category	Intended users	Process coverage in trend-radar	Relevant AI capabilities	Data and source model	Source traceability	Workflow automation	Human control options	Collaboration & communication	SME fit
ChatGPT	Generic AI assistant	Individual analysts, managers, teams, researchers	Brief development, keyword generation, source/material synthesis, coding support, trend drafting, scenario prompts, reporting	Generative AI; NLP; retrieval/research features where enabled	User-provided content; connected data and web retrieval depending on configuration	Configuration-dependent; source inspection must be actively required	Low in standard chat; medium to high where projects, connectors, automation or research modes are used	High at interaction level: users can prompt, revise, reject and regenerate outputs	Shared workspaces and export/sharing options depend on plan and configuration	High for initial analyst-led work; requires prompt, source and review discipline
Claude	Generic AI assistant	Individual analysts, managers, researchers, document-heavy teams	Research-brief design, document analysis, source/material synthesis, coding support, trend-profile drafting, narrative and critical review	Generative AI; NLP; long-context analysis; retrieval/research features where enabled	User-provided documents and connected/web sources depending on configuration	Configuration-dependent; it is not inherently a citation-first research environment	Low in standard use; medium where research, connected sources or integrations are available	High at interaction level: users control prompts, supplied evidence, revisions and approval	Sharing and team features depend on plan and configuration	High for analyst-led synthesis, especially where source collections must be read and compared
Gemini	Generic AI assistant	Individual analysts, business users, researchers, teams using Google services	Research-brief design, keyword generation, document/material synthesis, classification support, drafting and reporting	Generative AI; NLP; multimodal analysis; Google-grounded or research features where enabled	User-provided material; Google ecosystem and web retrieval depending on configuration	Configuration-dependent; grounding and citations must be checked in the specific mode used	Low in standard chat; medium to high in research, workspace and agentic configurations	High at interaction level; users review and revise outputs	Collaboration can be strong in Google Workspace environments	High where Google Workspace is already established; otherwise moderate
Perplexity	Source-oriented AI research tool	Individual analysts, researchers, managers and small teams	Topic exploration, horizon scanning, source discovery, comparison of viewpoints, preliminary	Retrieval-augmented generation; generative AI; NLP; multi-step research	Primarily web retrieval with AI synthesis; source access depends on	Relatively high for retrieved web material because responses normally present	Low to medium; research modes may perform multi-	High at interaction level; users choose queries, inspect sources and decide	Limited persistent team collaboration compared with foresight platforms	High for rapid, source-linked scanning; does not provide a full radar-management environment

Tool / framework	Tool category	Intended users	Process coverage in trend-radar	Relevant AI capabilities	Data and source model	Source traceability	Workflow automation	Human control options	Collaboration & communication	SME fit
			evidence synthesis and source recap	functions depending on mode	available web content and plan	source links; source quality still requires review	step web retrieval	which evidence to retain		
AgentGPT	Generic agent-building tool	Individual experimenters, analysts and technically confident users	Multi-step research planning, basic goal decomposition, preliminary scanning and structured summaries	Agentic workflows; generative AI; retrieval/tool use where configured	Depends on the available tools, APIs and source connections	Configuration-dependent; traceability depends on how retrieval and logging are implemented	Medium to high for defined multi-step goals	Human supervision is necessary because intermediate decisions and outputs require validation	Limited compared with dedicated enterprise or foresight platforms	Moderate to low for typical SMEs unless there is technical confidence and a narrow, well-defined use case
Lindy	No-code AI-agent and workflow-automation platform	Business users, operations teams and small teams seeking automation	Scheduled collection, monitoring, triage, routing, summaries, notifications and recurring reporting	Agentic workflows; generative AI; integrations; trigger-action automation	Connected applications, e-mail, calendars, spreadsheets, databases, RSS or other configured services	Dependent on connected source systems and workflow logging	High for recurring, trigger-based workflows	Users design the workflow and can introduce review steps and approvals	Can route outputs through shared business tools, e-mail or collaboration systems	Moderate to high where a clear repeatable process already exists; configuration work is required
CrewAI	Multi-agent development framework/platform	Developers, technical teams and organisations building custom agent workflows	Custom multi-stage processes: retrieval, research, filtering, clustering, critique, citation checking and reporting	Multi-agent workflows; generative AI; tool use; retrieval; memory and reasoning options	User-defined sources, models, tools, knowledge bases and integrations	Strong potential if source logging and citations are explicitly engineered; not automatic	High for custom orchestrated workflows	Potentially high but depends on workflow design, approval gates and interface implementation	Supports team and enterprise deployment, but requires development and operational set-up	Low for most individual SMEs; potentially useful through a service provider, association or technically capable partner
AutoGen	Multi-agent development framework	Developers, researchers and technical teams	Custom conversational or multi-agent workflows for research, critique, tool use, coding, analysis and reporting	Multi-agent workflows; generative AI; tool use; retrieval depending on implementation	User-defined models, tools, APIs, documents and data sources	Entirely implementation-dependent; traceability must be designed into the workflow	High for custom workflows, but requires technical development	Potentially high through designed review gates; not inherent in the framework	Depends on the application built around the framework	Low for typical SMEs without programming capacity; relevant for bespoke solutions or academic experimentation

Tool / framework	Tool category	Intended users	Process coverage in trend-radar	Relevant AI capabilities	Data and source model	Source traceability	Workflow automation	Human control options	Collaboration & communication	SME fit
IBM watsonx	Enterprise AI-development, data and governance platform	Enterprise IT, data-science, AI-governance and business teams	Development of tailored end-to-end systems: internal/external retrieval, classification, signal ranking, AI assistants, agents, analytics and governed deployment	Generative AI; retrieval; NLP; machine learning; agentic workflows; model/data governance	Enterprise data, user-provided knowledge bases and external sources, depending on implementation	Potentially high through enterprise data governance, logging and implementation design	High where systems are designed, integrated and deployed	High potential through role-based access, governance and workflow design	Enterprise-scale collaboration and integration potential	Low for most individual SMEs due to implementation and resource requirements; more plausible through shared-service models
Nextatlas	Purpose-built trend-forecasting and consumer-intelligence platform	Consumer-insight, brand, innovation and strategy teams	Detection and monitoring of consumer, cultural and market signals; trend exploration and consumer-oriented forecasting	Proprietary AI; pattern detection; forecasting; generative reporting; consumer intelligence	Proprietary early-adopter and digital-culture data, with possible integrations depending on plan	Relatively strong visibility into platform trend outputs; underlying proprietary data and methods are not fully transparent	Medium to high for always-on consumer-trend monitoring	Users can query, explore and interpret results; control over underlying data/model logic is limited	Platform reporting and stakeholder communication features	Moderate for SMEs with a consumer-innovation focus and available budget; less complete for regulatory, scientific or industrial scanning
Futures Platform	Purpose-built collaborative foresight platform	Foresight professionals, innovation teams, strategy teams and leadership groups	Horizon scanning, signal and trend management, prioritisation, scenario work, collaboration, foresight radar visualisation and reporting	Foresight-specific analytics and structured content; AI capabilities may vary by product configuration	Curated platform content, user-entered information and organisational material depending on implementation	Depends on the platform's source and content fields; source practices should be checked in the selected configuration	Medium to high for structured, repeated foresight workflows	Users can assess, discuss, prioritise and interpret content; human participation is central	High: collaboration, shared interpretation, reporting and radar views are core use cases	Moderate: likely more suitable for SMEs with an established foresight process or network/association support
FifthRow	Purpose-built AI-native research and foresight platform	Strategy, consulting, innovation and research teams	Automated or semi-automated horizon scanning, research synthesis, trend and uncertainty identification, scenario	Agentic workflows; generative AI; retrieval; synthesis; multi-step research	Provider-managed AI research workflows with source-supported outputs; source	Potentially source-supported, but actual traceability and auditability must	High for multi-step research and strategy workflows	Users define tasks and review outputs; degree of control depends on	Structured outputs for strategy and consulting work; collaboration features depend on offering	Moderate to low for individual SMEs unless pricing, transparency and workflow simplicity are favourable

Tool / framework	Tool category	Intended users	Process coverage in trend-radar	Relevant AI capabilities	Data and source model	Source traceability	Workflow automation	Human control options	Collaboration & communication	SME fit
			development and strategic implications		model depends on application configuration	be examined in use		workflow configuration		
FIBRES	Purpose-built trend-radar and foresight platform	Innovation, foresight and strategy teams; associations and collaborative networks	Organising data points, signals and trends; linking items; categorisation; radar visualisation; review; communication; maintenance and dissemination	Trend-data management; visualisation; potentially AI-supported analysis depending on configuration	User-entered, imported and platform-managed trend data; source model depends on the organisation's data practice	Potentially high if source links, dates, metadata and evidence fields are consistently retained	Medium: supports persistent radar management; automation depends on enabled integrations and features	High: users can add, edit, assess, link, approve and update radar entries	High: shared repository, team communication, radar publication and dissemination functions	Moderate to high for SMEs, particularly where a structured living radar is required and the platform cost is proportionate
Itonics	Purpose-built innovation, technology-intelligence and foresight platform	Innovation managers, technology scouts, foresight teams and enterprise innovation functions	Trend and technology scouting, evaluation, opportunity assessment, portfolio management, roadmapping, collaboration and innovation decision support	Foresight and innovation analytics; automation/AI capabilities vary by module and configuration	Platform-managed, user-entered and connected external/internal data, depending on configuration	Configuration-dependent; quality depends on retained source and evidence records	Medium to high for integrated innovation-management workflows	High potential through collaborative assessment, decisions and portfolio processes	High: collaborative evaluation, portfolio, roadmapping and reporting functions	Moderate to low for individual SMEs; potentially suitable for larger SMEs, associations or shared-service arrangements

Note: The table is a structured classification of documented tool characteristics and likely workflow roles at the time of the market scan. It does not provide an independent technical evaluation, benchmark, procurement recommendation or ranking of the listed tools. Product capabilities, data practices, citation behaviour, integrations, subscription conditions and governance features vary by version, plan, organisational configuration and date of access.

Jacobsen on Management