


nextFood

An AI product-innovation platform for food and beverage brands that continuously harvests global market signals, generates market-ready product concepts together with packaging visuals, and validates them against simulated target groups before a single sample is produced.



Tech Stack

Next.js • React • TypeScript • Tailwind CSS • shadcn/ui • Zustand • Node.js • BullMQ • Redis • Python • FastAPI • Scrapy • Supabase • PostgreSQL • pgvector • OpenAI • Google Gemini • Ideogram • fal.ai • Replicate • sentence-transformers • TinyTroupe • Zyte • Firecrawl • AWS SQS • Render • GitHub Actions • Cypress

The Challenge

Food and beverage innovation teams work against long odds, and the research meant to de-risk a launch is the slowest part of the process. Scouting the market means manually trawling retailer shelves, trade publications and analyst PDFs; turning that material into a written concept absorbs weeks of workshops before anything exists on paper.

General-purpose AI tools did not close the gap. They could produce copy, but they had no memory of what a brand had already launched, no view of what competitors put on shelf last month, no understanding of container formats, category rules or a specific market's consumer preferences, and no way to show what the finished pack would actually look like.

Consumer validation was the second bottleneck. Traditional panels are expensive and take weeks to organise, so teams had to commit to a direction long before they had any evidence for it. What was needed was one system that keeps watching the market, generates concepts that respect brand and market constraints, renders shelf-ready visuals, and tests the result against a target group in minutes rather than months.

Target Users

- Food and beverage innovation and NPD teams
- Brand and category managers at consumer goods companies
- Product development agencies and innovation consultancies
- Private-label and retail own-brand teams

Solution

nextFood turns product innovation into a single guided pipeline. Managed spiders and crawlers keep a brand's Data Silo stocked with structured, de-duplicated market intelligence from retailers, trade press and internal reports. Teams write a short briefing, pull in the trends and insights that matter, choose a visual direction, and set how far the AI should push beyond the current range. From there the platform generates complete product concepts with packaging visuals, has them rated by simulated target-group panels, and exports the winning idea as a branded one-pager or product video. Work that used to take months of workshops and panel research now happens inside one workspace.

Key Features

Managed Trend Harvesting: 25 production spiders pull products and articles from European, UK and US grocery retailers and trade publishers, on demand or on a cron schedule.

AI Trend Structuring: Web pages, product listings and PDF reports are normalised into one strict schema covering descriptions, key drivers, implications, ingredients and SEO keywords.

Semantic Duplicate Guard: Vector embeddings score every incoming trend against the existing library, so near-duplicates are flagged with their closest match before they reach a silo.

Hybrid Trend Search: Progressive keyword-plus-vector search across a brand's connected silos, with date and image filters, Google Trends curves and keyword search volume.

Brand-Aware Concept Generation: A reasoning model writes the complete concept — name, claim, ingredients, nutritional and flavour profile, USP, backstory and go-to-market plan, inside the brand's category, target group and market rules.

Packshot Generation and Editing: Packaging visuals are rendered in the brand's saved Look, then refined through instruction-driven tweaks, remixes, masked edits, text repair and upscaling.

Innovation Grade Control: One dial moves the AI from line extension to adjacent category innovation to disruptive breakthrough, mapped onto model reasoning effort and temperature.

Simulated Target-Group Testing: Up to 500 generated market-specific personas evaluate the concepts and return 1–5 ratings for overall appeal, ingredients and product name, plus a winner.

Multi-Market and Bilingual Output: Concepts honour per-country consumer context, with names and claims written in the market language and the remaining fields in the user's language.

Multi-Tenant Governance: Companies, brands, silos, looks and users are administered centrally with generation quotas, role-gated routes and per-project user assignment.

Workflow Diagram



Technical Architecture



Challenges & Solutions

Challenge 1: Generation That Outlives a Web Request

Challenge: A single concept takes minutes to produce because a reasoning model writes the copy and an image model renders the pack. No HTTP request survives that, and users cannot be asked to sit on a loading screen.

Solution: Every generation became a durable queue job. Three dedicated queues for ideation, packshots and brand assets are drained by a worker with five-minute locks, stall recovery and three retries with exponential backoff, while job rows in the database hold the state the UI polls.

Challenge 2: Every Retail Source Behaves Differently

Challenge: Grocery retailers hide catalogue data behind rendered pages, pagination quirks and bot protection, while trade publications are simple to read but wildly inconsistent in structure.

Solution: Ingestion was split into two tiers. Protected retail sites are handled by purpose-built Scrapy spiders running on managed Scrapy Cloud with AI-assisted product extraction and ban-aware retries; simpler sources go through on-demand crawling. Both paths land in the same normalised trend schema.

Challenge 3: Silos Filling Up With Near-Duplicate Trends

Challenge: The same product launch is covered by several publications and reappears on every scheduled run, so a silo quickly fills with variations of one story and buries the genuinely new signal.

Solution: Each record is embedded locally into a 384-dimension vector and matched against the whole library through a vector similarity function. Trends are stored with a similarity score and their nearest existing neighbour, and retrieval runs through a progressive hybrid search that combines full-text and vector ranking.

Challenge 4: Concurrent Image Jobs Overwriting Each Other

Challenge: Packshot and brand-asset jobs for the same idea finish independently. Reading an image array, appending to it and writing it back meant whichever job finished last silently discarded the other's output.

Solution: The appends were moved into atomic database functions, so each worker merges its results server-side in a single statement. Concurrent jobs now accumulate instead of competing, and a staleness flag on the idea tells the UI when visuals need regenerating.

Challenge 5: Multi-Provider AI Latency Was Invisible

Challenge: A single user action could fan out across several model providers, and when a flow felt slow there was no way to tell which call was responsible or whether a change had actually helped.

Solution: A shared flow tracker now times every step of every AI route. API responses carry performance metadata and background flows emit structured timing logs, which turned an optimisation pass into measurable work: independent provider calls were parallelised and keyword generation cached, cutting roughly eighteen seconds off the trends and insights flow on the measured runs.

Challenge 6: Consumer Testing Too Slow to Inform Early Decisions

Challenge: Teams needed a read on appeal while the concept was still cheap to change, but commissioning a real panel costs weeks and money that early-stage ideas do not justify.

Solution: A simulation service generates up to 500 personas constrained to the target market, with enforced diversity of age, profession, seniority and region. Each persona gives a first impression and then a detailed evaluation, from which structured 1–5 scores and a final pick are extracted and aggregated. Run locking and duplicate-request guards keep a single simulation authoritative.

Production & Scalability

The platform is built as a set of independently deployable services so that the web experience, the AI workloads and the scraping fleet can be scaled, released and failed independently of each other. Long-running work is queued rather than held open, every external integration is isolated behind a service boundary, and the same control plane drives both the development and production environments.

Service-Oriented Architecture

Five deployables cooperate over HTTP and a shared queue:

- Next.js web application and product API
- Node background worker consuming three job queues
- FastAPI AI service for trend extraction, embeddings and panel simulation
- FastAPI orchestration service for scraping, scheduling and batch ingestion
- Managed Scrapy Cloud fleet running 25 retailer and publisher spiders

Authentication & Authorization

Authentication is handled by managed sessions with password reset flows, layered with route-level guards that separate administrative consoles from the end-user workspace. Access is scoped by company, brand and project assignment, and every brand carries a generation quota that is enforced before work is queued.

Background Processing

Anything that cannot complete inside a request is queued and processed asynchronously:

- Product concept ideation and packshot rendering
- Additional packshot sets and brand asset generation
- Trend extraction, embedding and de-duplication
- Target-group simulation runs
- Scheduled spider execution and CSV batch ingestion

Progress & Notifications

Because generation runs for minutes, job state is persisted and surfaced rather than assumed. The interface tracks each job to completion or failure, and a notification layer queues toasts and marks the browser tab title and favicon when results arrive while the user is working elsewhere, so no one has to watch a spinner.

Security

The platform follows production security practices throughout:

- Session-based authentication with role-gated routes
- Company and brand level data isolation
- Secrets supplied entirely through environment configuration
- Time-limited signed URLs for all generated and uploaded assets
- Server-side validation on every mutating API route
- ORM and parameterised database access throughout

Database & Performance

Data is modelled for both transactional work and semantic retrieval:

- PostgreSQL holds roughly thirty tables spanning companies, brands, projects, ideas, silos, jobs and schedules
- pgvector stores trend embeddings behind progressive hybrid-search functions
- Object storage serves packshots, brand assets and scrape exports via signed URLs
- Redis backs the job queues and short-lived response caches
- Atomic database functions handle concurrent array appends and quota accounting

Observability

Every AI route and background flow is instrumented with step-level timings that are returned in API metadata and emitted as structured logs, and scraping jobs write an auditable event trail of status changes, item counts and failures. Performance work is therefore measured against recorded baselines rather than estimated.

Scalability

Each tier scales independently as demand grows.

Component	Scaling Strategy
Web application	Stateless Node service replicated behind a load balancer
Background workers	Additional worker processes per queue for higher generation throughput
AI service	Horizontally replicated instances with the embedding model loaded once per process
Scraping fleet	Managed Scrapy Cloud units scaled with the number of active spiders
Batch ingestion	Queue-buffered batches sized to record volume and drained in parallel
Redis	Managed Redis with per-environment queue namespacing
PostgreSQL	Managed database with vector indexes, backups and replication support

Deployment

The web application is deployed as a managed cloud service in the Frankfurt region, with a continuous integration pipeline that installs, builds and runs an end-to-end browser suite on every change to the main branch. Spiders are deployed to managed Scrapy Cloud, and every cross-service payload carries an explicit environment target so a single control plane can drive development and production without code changes.

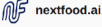
My Responsibilities

As the primary developer on the platform, I led the design, development, and delivery of the core system across all four repositories, contributing the majority of the codebase while collaborating with a small team on product design and individual scraping spiders. My responsibilities included:

- Designed and developed the nine-step innovation workspace in Next.js, covering briefing, discovery, market insights, visual trends, ideation, concept development, strategy, and testing.
- Built the asynchronous generation pipeline, including job queues, background workers, state management, retry logic, and failure handling.
- Implemented the AI orchestration layer, including prompt engineering, structured output schemas, provider selection, and reasoning strategy based on innovation complexity.
- Developed the FastAPI AI services for trend ingestion, web crawling, PDF parsing, document structuring, embedding generation, and semantic de-duplication.
- Built the scraping control plane, managed Scrapy Cloud integrations, retailer and publisher spiders, scheduled crawls, job tracking, and queue-based CSV processing.
- Designed the multi-tenant database schema and implemented hybrid search, atomic operations, quota management, and supporting database functions.
- Developed the AI target-group simulation engine, including persona generation, diversity constraints, multi-stage evaluation, structured scoring, and idempotent execution.
- Implemented multilingual support for German and English, including market-specific language resolution rules for AI-generated content.
- Optimized AI performance by instrumenting latency, profiling bottlenecks, and introducing parallel processing and caching strategies.
- Configured the CI/CD pipeline, end-to-end browser testing, and environment management across development and production deployments.

Gallery

Project Dashboard



Platform

BRIEFING

DISCOVER

- Market Insight
- Trends & Insights
- Visual Trends

DEVELOP

- Innovation Paths
- Creative Techniques

GENERATE

- Idea Room

TEST

- Target Group Testing

Functions

Save Project

Welcome Back, [anna.keller@feldkraft.demo](#)

Your Projects

+ Create New Project

Preview	Name	Created Date	Actions
	Plant-based protein range 2027	March 18, 2026	 
	Chilled snacking Q4	February 27, 2026	 
	Kids breakfast reformulation	February 11, 2026	 
	Private-label protein tender	January 22, 2026	 

AN [anna.keller@feldkraft.demo](#)

 Runs 34.00 / 120

Briefing Panel

nextfood.ai

Platform

BRIEFING

DISCOVER

DEVELOP

GENERATE

TEST

Functions

Save Project

Market Insight

Trends & Insights

Visual Trends

Innovation Paths

Creative Techniques

Idea Room

Target Group Testing

Brands

Objective

Objective Details

Category

Target Group

Container Form

Target Market

Ingredients

Features

Innovation Grade

Plant-based protein range 2027 | Briefing

Invite Users

English

CHALLENGE

The Briefing

Hint: Define your project foundation here. All information from the briefing is used throughout the workflow: in Trends & Insights generation, Visual Trends, Generation of Ideas, and Market Research.

Feldkraft (demo)

Plant-based protein range 2027

Extend the chilled range with three plant-protein RTDs for 2027, using EU-grown protein sources and no added sugar. Must work in a 330 ml can and hold a price point below the branded leaders.

Dairy alternatives

Gen-Z, health conscious, urban

330 ml can

Germany

oat protein, lupine, aronia

high protein, no added sugar, gut health

Choose how far the AI should push beyond your current products.

Line Extension / Renovation

Incremental improvements and safe updates close to existing products.

Adjacent / Category Innovation

Stretch into nearby categories with new twists while staying plausible for the brand.

Breakthrough / Disruptive Innovation

Bold, disruptive concepts that can redefine the category and surprise the market.

Start Generation

ANanna.keller@feldkraft.demo

Runs 34.00 / 120

Trends & Insights

Platform

BRIEFING

DISCOVER

- Market Insight
- Trends & Insights
- Visual Trends

DEVELOP

- Innovation Paths
- Creative Techniques

GENERATE

- Idea Room

TEST

- Target Group Testing

Functions

Save Project

Plant-based protein range 2027 | Trends & Insights

Invite Users

English

?

IMPULSES FOR IDEATION

Trends & Insights

Hint: Review and select trends that inspire your product ideas. Selected trends influence the ideation step and are referenced in the Details view.

Regional Protein Sourcing

Shoppers increasingly read the origin line before the nutrition panel. Lupine, faba bean and rapeseed protein grown inside the EU are moving from niche health stores into mainstream grocery ranges.

Key Drivers: Import-cost volatility, carbon labelling on shelf, and national campaigns promoting domestic crops.

Implications: Product Idea: Lead with a named growing region on front of pack. Marketing Idea: Publish the farm partner alongside the launch.

Functional Cold Brew Hybrids

Protein and caffeine are converging into one occasion. Ready-to-drink hybrids that replace both the breakfast shake and the takeaway coffee are the fastest-moving sub-segment in chilled.

Key Drivers: Commuter time pressure, decline of the sit-down breakfast, and the normalisation of functional additions in coffee.

Implications: Product Idea: A 330 ml can carrying both claims. Storytelling: Frame it as one fewer decision before 9 a.m.

Fermentation as a Clean-Label Signal

Fermented bases are being used to deliver acidity, shelf life and gut-health cues at once, letting brands drop additives from the declaration entirely.

Key Drivers: Consumer suspicion of long ingredient lists and rising awareness of the gut-brain axis in mainstream media.

Implications: Product Idea: Use a fermented fruit base instead of added acid. Marketing Idea: Make the live-culture count a hero number.

Client Insights

Low-Sugar Indulgence

The trade-off consumers refuse to make is taste. Dessert-adjacent flavour profiles with no added sugar consistently outperform austere health positioning in blind tests.

Key Drivers: Sugar-reduction targets, better sweetener systems, and fatigue with clinical health branding.

Implications: Product Idea: Hazelnut and cocoa profiles at zero added sugar. Storytelling: Sell the treat, footnote the nutrition.

AN

anna.keller@feldkraft.demo

Runs 34.00 / 120

Market Insight: Hybrid Trend Search

Platform

BRIEFING

DISCOVER

- Market Insight
- Trends & Insights
- Visual Trends

DEVELOP

- Innovation Paths
- Creative Techniques

GENERATE

- Idea Room

TEST

- Target Group Testing

Functions

- Save Project

Plant-based protein range 2027

Market Insight

Invite Users

English

GLOBAL PRODUCTS FOR INSPIRATION

Market Insight

plant protein rtd no added sugar

All

Last 90 Days

Hint: Select a few trends cards – they're used in Ideation and referenced in Details, helping you generate better ideas.

Oat Protein Drink, Chocolate 330 ml

Chilled oat protein RTD at 20 g protein, chocolate variant added to the core range in Q1.

Published: 2026-03-14

View Details

Retail listing - UK grocery

Lupine Barista Drink 1 l

Lupine-based barista drink positioned on regional sourcing and foam stability.

Published: 2026-03-02

View Details

Retail listing - NL grocery

Plant Protein Shake, Vanilla

Multi-source pea and faba blend, no added sugar, sold as a four-pack.

Published: 2026-02-21

View Details

Retail listing - FR grocery

Clear Vegan Protein, Berry

Clear-format vegan protein positioned against dairy whey on lightness and digestibility.

Published: 2026-02-09

View Details

Marketplace listing

Own-Brand Protein Drink, Coffee

Private-label entry undercutting branded RTDs by roughly 30 percent at parity protein.

Published: 2026-01-27

View Details

Retail listing - DE grocery

Fermented Oat Yoghurt Drink

Fermented oat base with live cultures, launched as a drinkable format for the breakfast occasion.

Published: 2026-01-15

View Details

Trade publication

AN

anna.keller@feldkraft.demo

Runs 34.00 / 120

Visual Trends & Looks

nextfood.ai

Platform

BRIEFING

DISCOVER

- Market Insight
- Trends & Insights
- Visual Trends

DEVELOP

- Innovation Paths
- Creative Techniques

GENERATE

- Idea Room

TEST

- Target Group Testing

Functions

- Save Project

Plant-based protein range 2027

Visual Trends

Invite Users

English

LOOK & FEEL

Visual Trends

Hint: Pick the art direction the packshot generator should follow. Your own design preset is combined with the selected looks.

Own Design Preset

Matte recycled board, single-colour print, ingredient photographed as-is on a raw surface. No gloss, no gradients.

168 / 400

Nordic Clean

Cool daylight, generous negative space, one accent colour.

studio still life, matte pack, pale grey seamless background, soft top light

Bold Grocery

High-contrast saturated colour built to win at two metres on shelf.

vivid gradient backdrop, hard key light, strong drop shadow, saturated pack

Craft Pantry

Kraft textures, warm neutrals, visible ingredient cues.

warm linen surface, scattered raw ingredients, window

Neo Retro

Seventies palette and rounded type for nostalgia-led ranges.

Clinical Functional

Clean white, precise geometry, laboratory-adjacent restraint.

Idea Room: Generated Concepts

nextfood.ai

Platform

BRIEFING

DISCOVER

DEVELOP

GENERATE

TEST

Functions

Save Project

Plant-based protein range 2027

Idea Room

Invite Users

English

IDEA ROOM

Overview of Product Ideas

Nordic Clean — v4

Feldkraft — functional core

Generate Ideas

Market Insight

Visual Trends

Creative Techniques

2 products selected

Nordic Clean + own preset

Copy Cat Strategy

Trends & Insights

Innovation Paths

Brand Concept

2 trends selected

Regional sourcing play

Feldkraft — functional core



Hafergold

Wach werden, ohne Kompromiss

Key Ingredients:

Oat protein isolate, green coffee extract, chicory fibre, sea salt

Target Flavor Profile:

Toasted oat with a light cocoa finish and a clean, dry aftertaste

Unique Selling Proposition:

Pairs 20 g of oat protein with a slow-release caffeine lift, so the morning shake and the morning coffee become one 330 ml can. Positioned for commuters who skip breakfast but refuse to skip protein.

Add to Target Group Testing



Lupinenkraft

Heimisches Protein, volle Power

Key Ingredients:

Sweet lupine protein, rapeseed oil, oat fibre, natural vanilla

Target Flavor Profile:

Mild and creamy with a nutty base and a faint hay-sweet top note

Unique Selling Proposition:

Swaps imported soy for lupine grown within 300 km of the dairy, answering the regional-sourcing question on the front of pack while holding a neutral taste profile that works in coffee.

Add to Target Group Testing

Concept Details & Packshot Editing

nextfood.ai

Platform

BRIEFING

DISCOVER

- Market Insight
- Trends & Insights
- Visual Trends

DEVELOP

- Innovation Paths
- Creative Techniques

GENERATE

- Idea Room

TEST

- Target Group Testing

Functions

- Save Project

Plant-based protein range 2027

Details

Hafergold

Wach werden, ohne Kompromiss

Pairs 20 g of oat protein with a slow-release caffeine lift, so the morning shake and the morning coffee become one 330 ml can. Positioned for commuters who skip breakfast but refuse to skip protein.

Key Ingredients
Oat protein isolate, green coffee extract, chicory fibre, sea salt

Target Flavor Profile
Toasted oat with a light cocoa finish and a clean, dry aftertaste

Marketing Strategy
Launch through chilled grab-and-go in urban stations before national grocery. Seed with commuter-route sampling at six stations, then convert to a subscription four-pack. Retail story leads on the named growing region, not the protein number.

Seed: 1284730193

More Options ▾

Fix Text

Tweak Packshot

New Packshots

Brand Assets

Remix Packshot

Video

Target Group Testing Results

Platform

BRIEFING

DISCOVER

- Market Insight
- Trends & Insights
- Visual Trends

DEVELOP

- Innovation Paths
- Creative Techniques

GENERATE

- Idea Room

TEST

- Target Group Testing

Functions

Save Project

Plant-based protein range 2027

Target Group Testing

Invite Users

English

Protein RTD concept screen — DE

PYTHON SIMULATION, GEN-Z, HEALTH CONSCIOUS, URBAN

120 participants · Germany · completed in 11 minutes

And the winner is...

Beerenbalance

Across 120 simulated participants drawn from the German market, Beerenbalance took the highest overall appeal at 4.6 out of 5 and won 47 of the final picks. Panellists consistently framed the fermented base as proof of a short ingredient list rather than as a health claim, and the tart profile was the only one described as refreshing rather than filling. Hafergold placed second on the strength of the coffee-and-protein occasion, but several participants could not explain what the name meant.

Beerenbalance

Hafergold

Lupinenkraft

Data Silo: Scraping Setup

nextfood.ai

Platform

BRIEFING

DISCOVER

DEVELOP

GENERATE

TEST

Functions

Save Project

AN

anna.keller@feldkraft.demo

Runs 34.00 / 120

English

?

SETTINGS

Scraping Setup

Start Now

Schedule

Select Spider

Aldi_Sued

Entry URL

https://www.aldi-sued.de/de/p/sortiment/kuehlung/milchalternativen.html

Select Frequency

Weekly

Pick a Date

March 2, 2026 - June 1, 2026

Time (UTC)

03:00

Start Scraping

Schedule Jobs

Sr#	Spider Name	Schedule (UTC Timezone)	Created Date	Status	Actions
1	Aldi_Sued	0 3 * * 1	2026-03-02	Active	
2	Albert_Heijn	0 4 * * 1,4	2026-02-24	Active	
3	Waitrose_New_Products	30 2 * * *	2026-02-18	Active	
4	Trending_Hunter	0 6 1 * *	2026-02-05	In-Active	
5	Tesco	0 5 * * 3	2026-01-29	Active	

25 spiders available in the connected Scrapy Cloud project.

Admin Console: Companies, Brands & Quotas

Platform

BRIEFING

DISCOVER

- Market Insight
- Trends & Insights
- Visual Trends

DEVELOP

- Innovation Paths
- Creative Techniques

GENERATE

- Idea Room

TEST

- Target Group Testing

Functions

- Save Project

English

?

Welcome Admin

All Companies

+ Create Company

Company Name	Runs used	Assigned brands	Assigned users	Actions
Feldkraft Group (demo)	34.00 / 120	3	7	
Nordmolke AG (demo)	88.00 / 150	5	12	
Hansa Retail Labs (demo)	12.00 / 60	2	4	

All Brands

+ Create New Brand

Preview	Name	Created Date	Actions
	Feldkraft (demo)	2026-01-08	
	Feldkraft Kids (demo)	2026-01-22	
	Nordmolke Fresh (demo)	2026-02-04	
	Hansa Own Label (demo)	2026-02-19	

Global Settings

Prompt Analysis Settings

Configure the analysis prompts used when generating base prompts from reference images in the Looks section.

Level 1 Analysis Prompt (Basic)

Describe the lighting, background and surface material of this packaging reference in one sentence. Do not name brands.

Level 2 Analysis Prompt (Detailed)

Return camera angle, focal impression, colour temperature, prop styling and negative instructions as separate clauses, suitable for direct use in an image prompt.

Save Changes

AN anna.keller@feldkraft.demo

Project Links

Live Link: www.nextfood.ai

Repositories: Can't share because of client privacy.

Final Takeaway

nextfood.ai shows what changes when AI is wired into a business process rather than bolted onto it. Market intelligence, concept generation, packaging visuals and consumer validation sit in one pipeline, each step feeding the next with structured data instead of copied text. The result is a production system that compresses a months-long innovation cycle into a working session, while keeping the brand rules, market context and human judgement that make a food or beverage launch credible.