

General Roadmap

Invite universal hands-on experimentation. Give everyone access to low-cost LLM tools such as ChatGPT, Mistral's Le Chat, or Google Gemini and encourage daily "sandbox" time to explore their capabilities.

Offer bite-sized training and inspiration. Point staff to short seminars, MOOCs, and curated success stories that demystify AI and spark ideas.

Collect and share emerging use-cases. Ask teams to document where AI helps, even modest time-savers, and circulate the best examples company-wide.

Identify internal AI champions early. Spot curious, motivated employees and give them extra resources or sandbox credits so they can mentor peers.

Protect calendar time for experimentation. Build regular "AI play hours" into team schedules so discovery is deliberate, not ad-hoc.

Create an AI Centre of Excellence (CoE). Form a cross-departmental forum with clear roles for knowledge-sharing, standards, and vendor scouting.

Articulate a concise AI vision. Describe in one page how AI supports organisational goals and the main problems it should solve first.

Publish a lightweight AI policy. Cover data privacy, intellectual-property, and security basics to give employees confidence to experiment responsibly.

Make AI a standing item on leadership agendas. Signal top-level commitment and allocate visible sponsorship.

Document and scale success. Turn emerging best practices into a living playbook—tools, prompts, and workflows mapped to business tasks.

Expand enablement beyond volunteers. Provide structured training paths so every employee, not just the keen few, can adopt AI tools.

Co-develop a focused AI strategy. Leverage the CoE and executive team to write a short, actionable plan that links AI investments to measurable business value.

Prioritise high-value pilot projects. Select 2-3 use-cases with clear ROI, defined success metrics, and executive sponsors.

Make AI literacy mandatory. Incorporate baseline AI competencies into onboarding, role profiles, and performance reviews.

Upgrade the policy to a full governance framework. Add ethical-use principles, security controls, and data-handling standards aligned with EU regulation.

Budget explicitly for AI. Allocate today's operating funds and forecast future investment for talent, tooling, and infrastructure.

Continuously monitor and refine the AI strategy. Use dashboards and periodic reviews to track adoption, ROI, and risk; iterate quarterly.

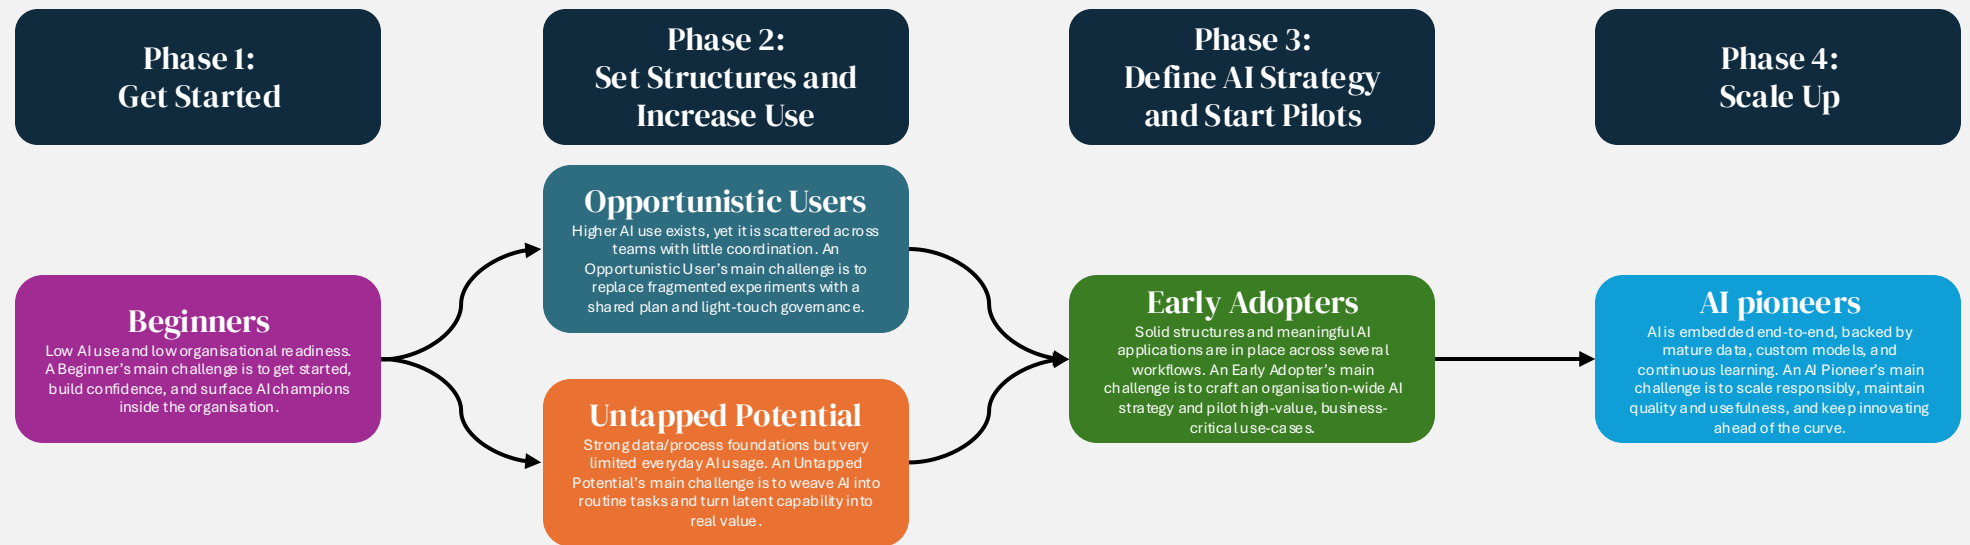
Choose the right build-vs-buy mix. Evaluate bespoke models, off-the-shelf solutions, and in-house development against cost, control, and speed.

Double-down on winners, sunset laggards. Scale successful pilots organisation-wide; retire or pivot projects that fail to deliver, documenting lessons learned.

Centralise data and guard quality. Conduct a data audit, eliminate silos, and establish enterprise data governance that underpins all AI work.

Evolve roles and keep the CoE vibrant. Introduce positions such as Prompt Engineer or Model Steward; celebrate innovation through internal showcases.

Maintain a culture of lifelong learning. Keep training catalogues fresh, provide advanced tracks, and recognise employees who push AI boundaries responsibly.



Research Roadmap

Run AI-assisted "literature-review labs"
Hold weekly sessions using ChatGPT or Mistral's Le Chat for summarising papers and research. Try out research specific tools like NotebookLM.

Pilot code-assistants for data cleaning and stats. Give volunteers GitHub Copilot/Cursor credits/subscriptions and capture productivity gains.

Launch a shared prompt repository. Use Slack or a Teams (or office suite of choice) channel where staff post successful prompts, datasets, and outcomes.

Nominate Research AI Champions. Empower 1–2 enthusiasts per unit to mentor peers and liaise with the CoE.

Schedule protected experimentation time. Block out weekly "play time" in project plans for safe experimentation.

Version-controlled prompt library.
Establish a "prompt & notebook library" for reuse and review. Make it easy to find and to update. Internal wikis work very well for this purpose.

Light data-governance checklist. Introduce naming conventions and an anonymisation checklist before AI processing. Update regularly when a new best-practice is found.

ROI working group. Form a cross-project team to track hours saved, value and insights gained from AI. Share those with the different teams, not just research if possible.

AI-enabled research templates. Make LLMs an integral part of survey design and sentiment analysis.

Advanced-technique training. Offer workshops on RAG, vector search and LLM-assisted coding. Over time, the technical vocabulary will improve.

Research-AI strategy. With the CoE and leadership, draft a roadmap that links AI to organisational KPIs.

High-value pilot selection. Launch 2-3 pilots (e.g. demand forecasting, large-scale web-scraping, semantic dataset search). It does not need to be complicated, but high value.

Compute & SaaS budget. Provide credits and specialised licences, with clear usage reporting to those who wish to go the extra mile.

Make AI a part of the onboarding process. New members of the research team should be shown from day one where AI is expected.

Ethics review upgrade. Expand the data-ethics process to include EU AI Act impact and bias checks.

Mandatory AI up-skilling. Make AI proficiency on par with other types of necessary and expected skills.

AI Insight Platform. Deploy an internal platform combining a knowledge graph, RAG chat and dashboards.

Build-vs-buy models. Decide whether to fine-tune a tourism-specific LLM or partner with external vendors.

Decide what privacy features are important to your organisation. Open source LLMs can be run locally and guarantee no data is used to train new models.

Continuous model monitoring. Automate bias, drift and data-quality alerts with clear ownership.

Open-source contributions. Publish non-sensitive tools or datasets to attract partners and talent.

AI Fellowship programme. Rotate staff through advanced projects to spread skills and avoid silos.

Collaborate with research institutions. Identify local universities interested in working together on AI projects.

Phase 1: Get Started

Phase 2: Set Structures and Increase Use

Phase 3: Define AI Strategy and Start Pilots

Phase 4: Scale Up

Beginners

Low AI use and low organisational readiness. A Beginner's main challenge is to get started, build confidence, and surface AI champions inside the organisation.

Opportunistic Users

Higher AI use exists, yet it is scattered across teams with little coordination. An Opportunistic User's main challenge is to replace fragmented experiments with a shared plan and light-touch governance.

Untapped Potential

Strong data/process foundations but very limited everyday AI usage. An Untapped Potential's main challenge is to weave AI into routine tasks and turn latent capability into real value.

Early Adopters

Solid structures and meaningful AI applications are in place across several workflows. An Early Adopter's main challenge is to craft an organisation-wide AI strategy and pilot high-value, business-critical use-cases.

AI pioneers

AI is embedded end-to-end, backed by mature data, custom models, and continuous learning. An AI Pioneer's main challenge is to scale responsibly, maintain quality and usefulness, and keep innovating ahead of the curve.

Marketing Roadmap

AI visual ideation. Encourage use of DALL-E 3, Midjourney, Runway or HeyGen to mock up concepts before studio work.

Run prompting workshops. Run guided sessions on prompting for ad copy, blog outlines and persona brainstorming.

Make a best-prompt gallery. Curate "best prompt - best output" examples in a digital space everyone has easy access to.

Two-week social challenge. Ask every marketer to create just one social post with AI, then compare results in a group call – learning by doing.

Set up regular inspiration sessions. Set aside some time every month to review what other marketing teams around the world have achieved with AI.

Organise weekly creative sprints. Test AI on low-risk assets and share learnings every Friday.

AI brand guardian. Assign a team member to vet tone-of-voice, inclusivity and compliance of AI outputs.

Have a brand-voice prompt library. Maintain do's & don'ts for each channel as a living document, preferably as an internal wiki.

Prompt-approval workflow. Add a simple review step so brand-critical prompts get a second pair of eyes before use at scale.

Quick-request form. Set up a Form so any team can ask the AI CoE to review or tweak AI content. This keeps oversight lightweight but accessible.

Policy quick-reference card. Create a marketing specific short AI policy in a one-pager format with "3 do's, 3 don'ts, and who to ask."

Send out a periodic AI roundup e-mail. Recap new tools, best prompts and any pitfalls spotted. This keeps everyone aligned with minimal fuss.

Use AI in more tasks than previously. Use LLMs to feedback and review more things: images, websites. Integrate them into more workflows, even it may not always seem like they fit in.

Use the collaboration features. Tools like ChatGPT can create projects that can be shared in teams, with specific overarching instructions and documents.

Try out more advanced tools. Platforms like Lovable or Replit allow users to build webpages/apps, which can result in faster prototypes and ideas testing.

Train AI agents with your own data. Feed internal manuals, guides or newsletters as training material to agents/custom GPTs so AI answers match your specific tasks.

High-value pilot selection. Launch 2-3 pilots (e.g. event gathering guide for tourists). It does not need to be complicated, but to focus delivering high value.

Monthly demo day. Invite teams to show one AI-created campaign asset and its performance. This should build confidence and a little internal FOMO!

Scale your successful AI pilots. Evaluate what worked and what didn't, rethink if there are any functions that can be transferred to other of your team's use cases or even other teams.

Make an annual tool benchmark. Once a year, test your chosen AI tool against two competitors on quality, cost and speed to keep pricing and performance on par with what exists on the market.

Keep a continuous-improvement board. Kanban column "AI Fixes & Ideas" lets anyone log issues or upgrades; review at the monthly CoE meeting.

Manage a AI knowledge base. Build a simple internal wiki with FAQs, short videos and compiled best-practices so newcomers can self-serve answers.

Maintain a risk register. Keep a living table of any AI-related issues (e.g., wrong facts, copyright flags) and how you fixed them. It helps show responsible governance.