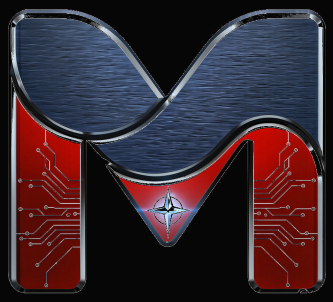

MICKAI EBOOK SERIES · PLAYBOOK No. 10

How Using AI Can Save Your Business Time.

A practical playbook for the UK SME owner and the regulated-enterprise CTO who wants to know what AI actually does to throughput, where it does not, and what the right adoption sequence is in 2026.

AUTHOR

Micky Irons

Founder and named inventor, Mickai LTD.

Crunchbase · LinkedIn · GitHub · mickai.co.uk

DATE · 18 May 2026 · v1

EBOOK · No. 10 IN A SERIES OF 14

Mickai LTD · Companies House 17166618 · press@mickai.co.uk · mickai.co.uk

UK IPO patent family GB2607309.8 to GB2610422.4 · Trade mark UK00004373277

TABLE OF CONTENTS

Contents

FOREWORD

A note from the author

AI is now a productivity instrument that competing UK SMEs are already using. This ebook is the practical playbook on what AI actually saves time on (document drafting, customer triage, CAD review, code refactor, transaction monitoring), what it does not, and the right adoption sequence for an owner-led business or a regulated-enterprise function in 2026.

It is written for the SME owner who wants to act this quarter and the regulated-enterprise CTO who wants to brief their executive team next month.

The Mickai substrate primitives are filed at the UK IPO across the GB2607309.8 to GB2610422.4 patent family. The trade mark Mickai is registered at UK00004373277.

Micky Irons

Founder and named inventor, Mickai LTD · 18 May 2026

PART I · THE THROUGHPUT QUESTION

Where AI lifts throughput, where it does not, how to measure

1. Where AI lifts throughput, by department

Document-heavy departments (legal, finance, regulatory, marketing communications) see the largest throughput lift from AI. Drafting, summarising, reviewing, and citation-checking are the workflows where the lift is consistently five to ten times. Customer-facing departments (sales, support, account management) see lifts of two to five times on triage, classification, and response drafting. Engineering departments see two to four times on code refactor, CAD review, and document classification.

Operations departments see uneven lifts; transaction monitoring and exception triage benefit, but supervisory and audit functions need careful design to avoid replacing one set of risks with another.

2. Where AI does not lift throughput

Tasks that depend on tacit organisational knowledge (a partner's judgment on a multi-party negotiation, a CTO's call on a strategic architecture decision, a CEO's read of a regulator's mood) do not benefit from current-generation AI. The model is not the bottleneck; the context is. Putting AI on these tasks produces plausible but low-trust output that is worse than the human alternative for the same time investment.

3. How to measure the lift in operator language

Throughput lift is measured per workflow, not per department. Pick three workflows, baseline them in human-hours, deploy AI assistance, remeasure in human-hours. The delta is the lift. The lift is then translated into the operator's own units (briefs drafted per week, claims processed per day, code merge requests cleared per sprint).

The substrate captures the workflow audit trail; the operator can verify the lift independently of the AI vendor's reporting.

PART II · THE ADOPTION SEQUENCE

From inventory to pilot to scale

4. Workflow inventory

Step one is to inventory the workflows. List every recurring process across the organisation, classified by department, frequency, and time investment. The top quartile by time investment is the candidate set for AI assistance. The bottom three quartiles are typically not worth automating.

5. Pilot selection

Step two is to pick the pilot. The pilot workflow should be high-volume, document-heavy, low data-sensitivity, and reversible. A common first pilot is internal document drafting (memos, policies, briefings). The pilot runs for a sprint, captures human-hour before and after, and produces a transferable artefact.

6. The throughput ledger

Step three is to keep a throughput ledger. Every AI-assisted workflow records its human-hour delta and its quality assessment in the operator's own measurement system. The ledger is the substrate-backed record of the AI's contribution to the business and is the artefact the board reviews quarterly.

PART III · THE STACK

What AI stack is right for your data class

7. Cloud AI, edge AI, on-device AI

Three stack options. Cloud AI (OpenAI, Anthropic, Google, AWS) is fastest to deploy, lowest upfront cost, highest data-egress risk. Edge AI (vendor hardware on premise) is moderate-cost, moderate-friction, retains data on operator network. On-device AI (operator-owned hardware running open or proprietary models) is highest-control, highest upfront cost, lowest data-egress risk.

8. The data-class boundary

The stack decision is made per data class. Low-sensitivity data (marketing copy, public-facing communications, code that is open-source-bound) can sit on cloud AI. Medium-sensitivity data (internal policy, draft contracts, customer correspondence) sits on edge AI or on-device AI. High-sensitivity data (regulated artefacts, trade secrets, personal data of customers under UK GDPR) sits on on-device AI under operator key custody.

9. The substrate position

Across all three stack options, the substrate question is whether the audit chain is held under operator key custody. The Mickai OAR primitive is stack-neutral; the same chain format wraps cloud AI calls, edge AI calls, and on-device AI calls. The substrate decision is independent of the stack decision.

PART IV · THE TWELVE-MONTH PLAYBOOK

Quarter-by-quarter milestones

10. Quarter 1, 2, 3, 4 milestones

Quarter 1

Workflow inventory, pilot selection, baseline measurement, first AI tool deployed on a single workflow with measured before and after.

Quarter 2

Three workflows running AI assistance, throughput ledger established, first quarterly board review of the ledger.

Quarter 3

First high-sensitivity workflow migrated to on-device AI under operator key custody, OAR chain in operation, browser-resident verifier deployed.

Quarter 4

AI assistance integrated into standard operating procedure across the top quartile of workflows, governance recital produced, twelve-month throughput lift quantified.

11. The CTO/CFO conversation

Twelve months in, the CTO and CFO sit down with the board and walk the throughput ledger. The conversation is about whether the realised lift justifies the spend, whether the substrate position is defensible, and whether the next twelve months expand the surface or consolidate the gains.

12. Closing

AI saves time on the right workflows under the right stack with the right substrate. Picking the wrong workflow, the wrong stack, or skipping the substrate produces the opposite of the intended outcome.

Engineering and finance leadership at any UK SME or regulated enterprise is open to a thirty-minute adoption-sequence briefing at any time. press@mickai.co.uk.

APPENDIX · ABOUT THE AUTHOR

Micky Irons

Founder of Mickai LTD (Companies House 17166618, England and Wales, registered office 20 Wenlock Road, London, N1 7GU). Named inventor on the Mickai SIOS patent corpus, recorded on the UK Intellectual Property Office public register at numbers GB2607309.8 to GB2610422.4. Trade mark Mickai registered at UK00004373277 (classes 9 and 42, filed 15 April 2026).

Before founding Mickai, Micky was a Sellafield site worker. The egress constraint observed from inside the regulated workstation is the engineering origin of the substrate described across the Mickai ebook series.

Profiles and links

mickai.co.uk · the canonical Mickai site.

crunchbase.com/person/micky-irons · founder profile.

linkedin.com/in/mickyirons · personal LinkedIn.

github.com/Micky-CMO · open-source position.

linkedin.com/company/mickai · Mickai LTD company page.

crunchbase.com/organization/mickyirons · Mickai LTD Crunchbase entry.

Email: press@mickai.co.uk

Colophon

Set in Inter Tight (Variable) and Inter Black. Brand voice audited under the Mickai AMT preflight gate; zero violations at publish.

References and further reading

· DORA-style engineering productivity studies, 2025 to 2026 series.

· Information Commissioner's Office, AI and data protection guidance.

· PRA Supervisory Statement SS1/23, model risk management for banks.

· Mickai brain taxonomy: mickai.co.uk/brains.

· Mickai trade mark UK00004373277, classes 9 and 42, filed 15 April 2026.