



Luís Rodrigues

Helping Leaders Turn AI into ROI | CPTO | Leading Digital Transformation Across FS, Telco & Government | Follow for posts on AI & business

58k followers • [View Profile](#)

8

POSTS / WEEK

1 PM, UTC

MAIN POSTING TIME

610

AVG REACTIONS
190 comments

98%

CALL-TO-ACTION
FREQUENCY

#1 SUCCESS FACTOR

He wins by turning scary, complex AI topics into simple “one-page” mental models leaders can use immediately. Those clear frameworks (usually a stack, checklist, or comparison chart) match his profile promise of “AI into ROI,” so people trust him, share the visuals, and then take the obvious next step: subscribe to the newsletter for the deeper playbooks.

FORMAT BREAKDOWN




Image	94%	650 avg reactions
Carousel	4%	380 avg reactions
Video	2%	420 avg reactions

7 Tactics You Can Steal

01 [Start by saying what most people get wrong, then fix it](#)

02 [Teach with a numbered stack and one picture people can save](#)

03 Split the topic into "not this" and "this" so it's clear

04 Talk about what breaks in real life, not just the shiny parts

05 Make "free learning packs" with big names, links, and a tiny question

06 End with one simple question that makes people answer about themselves

07 Ask people to comment a word before you send the free resource

7 Tactics You Can Steal

1 Start by saying what most people get wrong, then fix it

He often opens with a blunt correction ("most people think..." or "stop doing...") and then replaces the bad idea with a clear set of rules. This works because people can't resist checking whether they're making the same mistake, and the "fixed version" is easy to share with a team. In his highest performers, this "you're thinking about it wrong" opening shows up in 8 of the 10 best posts, including some above 3x relative engagement. To copy it: write the common wrong belief in quotes, say it's wrong in one short line, then teach your 3–6-point "better way."

[Example 1 →](#)

[Example 2 →](#)

2 Teach with a numbered stack and one picture people can save

He turns the idea into a simple numbered "stack" so readers can skim it fast, then backs it with one reference image that works like a cheat sheet. This works because the post becomes something people bookmark, screenshot, and repost to explain it to others. The "AI is a stack/layers" series averages 1.97x relative engagement across 8 posts, and image posts are his best-performing format overall (1.26x relative engagement across 93 posts). To copy it: pick one metaphor (stack, ladder, layers), write 5–7 numbered lines with plain-language labels, and attach one clean diagram that matches the list.

[Example 1 →](#)

[Example 2 →](#)

3 Split the topic into "not this" and "this" so it's clear

He draws hard lines between look-alikes (what something is not) and the real thing (what it is), then gives a short checklist of observable traits. This works because it stops fuzzy debates and gives readers words they can use in meetings the same day. Several of his best posts use this boundary-setting style, including ones above 2.5x–3x relative engagement. To copy it: list 3 common "not this" examples, explain in one line why each fails, then give a 4–6-point checklist of what the real version must do.

[Example 1 →](#)

[Example 2 →](#)

4**Talk about what breaks in real life, not just the shiny parts**

He doesn't stop at "cool demo" talk; he lists the messy stuff that decides whether something works at work (rules, logs, cost limits, safety checks, and how you undo mistakes). This works because it signals real experience and gives leaders a way to judge risk, not just excitement. The posts that lean hard on these "real-world failure points" reach roughly 2.7x to 3.8x relative engagement in his top set. To copy it: add a section called "What goes wrong" and list 5–8 failure points plus the simple fix for each.

[Example 1 →](#)[Example 2 →](#)**5****Make "free learning packs" with big names, links, and a tiny question**

He regularly turns public learning into a clean pack: one strong "this is free" headline, a numbered list of links, and one easy question at the end. This works because it saves time, and the big-name sources act like instant trust stickers. Across his broader posting, big-brand free course/curriculum posts average 1.39x relative engagement across 11 posts, and one of his top performers is exactly this kind of list. To copy it: pick one trusted source, list 7–15 links with one-line descriptions, then ask "Which one will you start with?"

[Example 1 →](#)[Example 2 →](#)**6****End with one simple question that makes people answer about themselves**

He often finishes by asking a question that helps the reader label themselves ("where are you spending time?" or "which one will you start?"). This works because it's easier to answer about yourself than to "add thoughts," so more people comment. Engagement questions show up in 61 posts, making it one of his most-used interaction tools. To copy it: avoid big open-ended questions, and instead ask for a choice, a starting point, or a quick self-rating.

[Example 1 →](#)[Example 2 →](#)**7****Ask people to comment a word before you send the free resource**

When he has a concrete asset (a one-page file, a PDF, a toolkit), he asks people to comment a single keyword or repost, and then he sends it. This works for two reasons: the comments boost reach fast, and he gets a list of people who clearly want the thing so he can message them. These direct-response posts average 1.45x relative engagement across 3 posts, and one generated 928 comments. To copy it: teach the idea first, show what's inside the resource, then ask for a one-word comment and actually deliver it quickly.

[Example 1 →](#)[Example 2 →](#)[Example 3 →](#)

Profile Breakdown

HEADLINE FORMULA

Helping Leaders Turn AI into ROI | CPTO | Leading Digital Transformation Across FS, Telco & Government | Follow for posts on AI & business

Helping [TARGET ROLE] achieve [BUSINESS OUTCOME] | [SENIOR OPERATOR TITLE] | Worked across [COMPLEX INDUSTRIES] | Follow for [TOPIC PROMISE]. Works because it leads with the result (ROI/impact), then proves "I've done this at scale" with a senior role plus regulated-industry scope, and finishes with a clear reason to follow.

BANNER STRATEGY

Doing AI before
it was trendy! ↘

Stop piloting AI
Start delivering impact

Follow me for emerging strategies in
AI-native product development

Big, high-contrast text that repeats the same promise as the headline/about (stop pilots, deliver impact) + 2-3 highlighted keywords + a directional arrow pointing down at the profile/CTA. Works because it turns the top of the profile into a fast "billboard": one clear outcome message, instantly scannable, with your eye guided to the next action.

Content Strategy

PILLAR	%	WHAT CRUSHES	EXAMPLES
AI systems explainers with simple pictures (stacks, definitions, architecture)	36%	Diagram-led explainers that turn the topic into a simple "stack/layers" mental model are the standout: the "AI is a stack/layers" series averages 1.97x relative engagement across 8 posts (and sits above the pillar's 1.57x average). What consistently underperforms here is deep infrastructure/news without a simple decision lens: 2 posts average 0.38x relative engagement.	Post 1 Post 2 Post 3
Curated learning resources people can use right now (free courses, curricula, repo lists)	26%	Big-name, free "course/curriculum" posts are a reliable engine: 11 posts average 1.39x relative engagement (above this pillar's 1.22x average). The format that wins is a tightly structured list (numbered items + who it's for + direct links). What tends to flop is repeat/low-novelty resource reposts: 3 posts average 0.50x relative engagement.	Post 1 Post 2 Post 3
Practical tools and workflows (prompts, coding setups, step-by-step how-tos)	18%	The strongest performers in this pillar are "here's the exact artifact/workflow" posts that people can copy and apply, especially when paired with a direct-response resource handoff (comment a keyword). Those direct-response lead-magnet posts average 1.45x relative engagement across 3 posts. Simple tool announcements without a clear "why it matters and how to use it" angle lag: 4 posts average 0.58x relative engagement.	Post 1 Post 2 Post 3
Leadership and org change (what leaders should do, hire, measure, and stop doing)	19%	This pillar performs best when the advice is specific and action-based (what to change in the team, the operating model, or the skill plan): 5 tactical, data-anchored posts average 1.08x relative engagement (above the pillar's 0.76x average). Abstract reflections and broad metaphors tend to sink: 3 posts average 0.43x relative engagement.	Post 1 Post 2 Post 3

Conversion Strategy

Their Offer(s)

Build What Matters newsletter (free subscription)

DM-based resource delivery (comment a keyword / repost to receive a PDF or link)

AI execution help via DM (consulting-style inbound)

Third-party product promotion: ChatLLM by Abacus AI (\$10/month)

How They Promote in Posts

He almost always teaches first and places the CTA at the end, so the post still feels complete even if you never click. Newsletter CTAs appear in 97 of 99 posts (97.98%), and the dataset is so CTA-heavy that it's not possible to cleanly compare "CTA vs no CTA" (there's only one clear no-CTA example). Higher-friction CTAs (comment a keyword, repost to receive, DM) are used much more selectively and usually tied to a specific deliverable, which helps keep trust high while still capturing high-intent leads.

[Example 1 →](#)

[Example 2 →](#)

[Example 3 →](#)


How They Promote in Profile

The profile pushes one main action everywhere: the headline and banner repeat the "AI into ROI / stop piloting" promise, the About section backs it with scannable proof blocks and ends with a follow CTA, and the Featured section is a single, clear newsletter subscription button—so high-intent visitors have an obvious next step.

[View Profile →](#)

Top 10 Posts

1



AI isn't magic.

He explains AI as a simple stack of layers and walks through a six-step progression from older rule-based systems to modern systems that can plan and act. The post is written like a short lesson, with a clear sequence and a visual that matches the explanation. It ends with a reflective question and a newsletter invite.

♡ 2720 💬 245 ↻ 274

2




Most people think Agentic AI is just "ChatGPT + tools"

He argues that real "agent-like" systems are not just a chat tool connected to apps. He lays out a layered model and focuses on the system pieces that decide whether it works in the real world, like safety checks, tracking, memory rules, and what happens when something fails. He closes by changing the evaluation question from "which tool?" to "what happens when it breaks?"

♡ 2676 💬 256 ↻ 290

3



AI isn't magic.

He repeats the "AI is a stack" idea and explains each layer in plain language, showing how each step builds on the one below it. The post uses short sections with bold labels so it reads well on mobile. It finishes with a question that prompts readers to share where they focus today.

♡ 2514 💬 306 ↻ 318

4



Stop calling everything an AI Agent.

He pushes back on the overuse of the word "agent" and explains how that confusion hurts real business results. He separates common categories (like chat interfaces, basic automation, and "search then answer" systems) from systems that can truly plan and act. He ends with a crisp contrast and a question about what readers are actually building.

♡ 2251 💬 297 ↻ 204

5



Most people think Agentic AI is just "ChatGPT + tools"

He repeats the core argument that "agentic" work requires more than a chat model plus tools. He breaks the system into layers and highlights the overlooked parts that cause failure in practice, like safety rules, tracking, memory policies, and recovery plans. He ends with a question that reframes how to judge solutions.

♡ 1803 💬 294 ↻ 164

6



Harvard released a full course on AI & prompting online.

He shares a free course from a trusted institution and turns it into a simple learning pack with direct links to sessions. Each item is described briefly so readers can choose what matters to them. He adds bonus resources and ends by asking which one people will start with.

♡ 1725 💬 278 ↻ 227

7



Stop calling everything an AI Agent.

He again clarifies what an “agent” is and isn’t, focusing on observable behaviors like planning, taking actions through real tools, remembering, and coordinating with other helpers. He uses everyday business examples (like systems touching customer tools and workflows) to keep it accessible. The post ends with a clear contrast and a question.

♡ 1815 💬 235 ↻ 132

8



This list saves you 100+ hours of research:

He promises a big time save and delivers a curated list of educational AI resources on GitHub, including star counts and direct links. He explains the surprising pattern he found while researching and shares how he filtered the data. He offers the script used so others can repeat the work.

♡ 1767 💬 197 ↻ 216

9



AI isn't magic.

He reframes AI as an architecture choice and compares a few system designs with clear trade-offs, including where they fail and what they're good for. He warns about common issues like systems making things up when they aren't grounded in real data. He ends by recommending an integrated approach that combines the strengths of each design.

♡ 1556 💬 279 ↻ 143

10



AI isn't magic.

He explains AI as a layered stack using a more conversational, less intimidating tone while still keeping the structure clear. He walks through the progression from older approaches to modern systems that can remember context, make plans, and use tools. He ends with a prompt that invites readers to share which layer they work on most.

♡ 1458 💬 256 ↻ 147