

Luís Rodrigues



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

49k followers • [View Profile](#)

8

POSTS / WEEK

1 PM, UTC

MAIN POSTING TIME

530

AVG REACTIONS
190 comments

99%

CALL-TO-ACTION
FREQUENCY

#1 SUCCESS FACTOR

He wins by turning confusing AI hype into simple, reusable “boardroom tools” (clear diagrams, clean categories, and real-world checklists) that make leaders feel confident making decisions. That same execution-first “AI into ROI” positioning is repeated everywhere on the profile, so the content earns trust fast—and the next step is always obvious (subscribe/follow), which steadily converts attention into an owned audience.

FORMAT BREAKDOWN

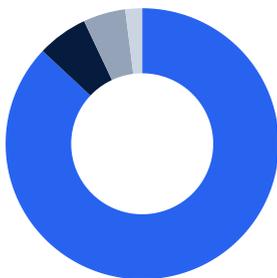


Image	87%	560 avg reactions
Video	6%	270 avg reactions
Carousel	5%	460 avg reactions
Text Only	2%	90 avg reactions

7 Tactics You Can Steal

01 [Turn big ideas into a simple stack people can screenshot](#)

02 [Start by correcting a common mistake, then give a clear definition](#)

03 Teach by comparing three options and saying what breaks in each

04 Hide a "real life" checklist inside the explainer

05 Make resource posts easy to save: short list plus one-line notes

06 Share a surprising business opinion, then prove it with simple math

07 Post like a TV show: same time, most days, with images

7 Tactics You Can Steal

1 Turn big ideas into a simple stack people can screenshot

He regularly takes a confusing topic and turns it into a short, numbered “stack” with a matching image. It works because people don’t just learn it—they can reuse the picture in meetings, so they save and share it. This is the most proven pattern: stack/diagram explainers average about 1.98x engagement across 20 posts, and two of the highest performers are over 4x. To copy it: name your layers, keep each layer to one line, and make the image match the exact same order and words.

[Example 1 →](#)

[Example 2 →](#)

2 Start by correcting a common mistake, then give a clear definition

Several top posts start with a firm correction (“stop calling X by the wrong name”) and then replace confusion with a simple checklist. People engage because it gives them words to push back at work, without needing to be an expert. This “definition cleanup” shows up in multiple of the highest-performing posts and reliably lands around 2.5x–3.6x relative engagement. To copy it: pick one misused term in your field, write a one-sentence correction, then list 3–5 “it counts only if...” rules.

[Example 1 →](#)

[Example 2 →](#)

3 Teach by comparing three options and saying what breaks in each

He often explains choices by lining up a few options side by side, then naming the tradeoff and the “gotcha” for each one. This works because readers can quickly spot where they are today and what they’re missing, which makes the post feel like a decision tool, not a lesson. One of his top posts using this approach lands around 2.6x relative engagement, and this “comparison chart” style repeats across his best-performing explainers. To copy it: choose three common approaches in your niche, write one benefit and one failure for each, then end with what the best version combines.

[Example 1 →](#)

[Example 2 →](#)

[Example 3 →](#)

4

Hide a “real life” checklist inside the explainer

His highest performers don't stop at explaining the idea—they include a practical list of what fails in the real world (things like guardrails, logs, rollback, costs, and coordination). That's why leaders and builders both share it: it helps them review projects and avoid embarrassing surprises. Two top posts that do this land around 3.0x–4.2x relative engagement, which is far above the baseline. To copy it: after your framework, add a “Before you ship this, check:” list with 6–10 items that prevent failure.

[Example 1 →](#)[Example 2 →](#)[Example 3 →](#)**5**

Make resource posts easy to save: short list plus one-line notes

When he shares learning resources, he packages them as a clean, numbered list and adds one quick “why this matters” line for each item. This works because it removes the hard part (searching and deciding) and turns the post into something people bookmark. Resource curation is a major pillar (28 posts averaging about 1.20x overall), and the best-curated sets climb much higher (a top example is near 3x). To copy it: pick one trusted source, list the parts people should start with, and end with a simple choice question (“Which one will you start?”).

[Example 1 →](#)[Example 2 →](#)[Example 3 →](#)**6**

Share a surprising business opinion, then prove it with simple math

One of his biggest non-technical hits starts with a bold claim, then walks through the numbers step by step until the conclusion feels unavoidable. People comment because they can argue with an opinion, but it's hard to argue with a clear chain of math. This matches what works in his leadership pillar: data-backed contrarian takes average about 1.80x, far above generic promo posts at about 0.28x. To copy it: make one strong claim, show 3–6 simple numbers with sources or assumptions, then end with one practical decision rule.

[Example 1 →](#)[Example 2 →](#)[Example 3 →](#)**7**

Post like a TV show: same time, most days, with images

He posts at a very steady pace (about 8 times per week) and almost always at the same time (1 PM UTC accounts for 94% of posts). That consistency trains your audience when to expect you, and it also makes it easier for you to keep the habit. He also leans hard into static images, which perform far better here (median likes: image 564 vs video 269 vs text-only 86). To copy it: pick one daily time you can hit for a month, prioritize one strong image per post, and only add video when you have a clear reason.

[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 DECISION-MAKER] turn [TREND/TECH] into [BUSINESS RESULT] | [SENIOR OPERATOR TITLE] | [PROOF BY CONTEXT: industries/scale] | Follow for [CONTENT PROMISE].
Why it works: it leads with a board-level outcome (ROI), immediately backs it with authority (CPTO) and credible environments (FS/Telco/Gov), then sets a clear expectation for what you'll learn if you 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 (stop experimenting → start delivering measurable impact) + highlighted keywords for fast skimming + a simple follow prompt.
Why it works: it turns the top of the profile into a one-sentence positioning statement that matches the feed content (practical execution), so visitors quickly understand “what you stand for” and what to do next.

Content Strategy

PILLAR	%	WHAT CRUSHES	EXAMPLES
Agentic AI & AI architecture explained with simple visuals	31%	Diagram-led “stack/layers/definitions” explainers are the strongest, most repeatable winners: this style shows up in 20 posts averaging about 1.98x relative engagement, and it dominates the highest-performing content. What underperforms inside this pillar is narrow technical optimization news without a manager-level frame (3 posts averaging about 0.42x).	Post 1 Post 2 Post 3
Curated learning resources people can save and come back to	28%	Clearly structured “free course/library/blueprint” curation is a reliable engine: the best cluster (7 posts) averages about 1.48x because it’s easy to bookmark and share. Performance drops when it feels like a generic resource dump or repeated list (4 posts averaging about 0.48x), especially if it’s text-heavy and not tightly packaged.	Post 1 Post 2 Post 3
AI tools and workflows, framed as decisions (what to use, when)	14%	Tool/workflow posts do best when they take a clear stance and help people choose (6 posts averaging about 1.29x). What drags results down is anything that reads mainly like a promo or pricing push (3 posts averaging about 0.54x), even if the tool is good.	Post 1 Post 2 Post 3
AI leadership, economics, governance, and “what changes next”	27%	The breakout performers here are data-backed, slightly contrarian business takes (5 posts averaging about 1.80x) because they give leaders a fresh lens and language for decisions. The weakest cluster is generic motivation, announcements, or overtly promotional updates (4 posts averaging about 0.28x).	Post 1 Post 2 Post 3

Conversion Strategy

Their Offer(s)

Build What Matters — free newsletter subscription

DM-delivered assets/frameworks (keyword DMs and “I’ll send you the guide/PDF” mechanics)

Affiliate/partner tool recommendations and discounts

Third-party course/certification promotions

Occasional hiring/role promotion

How They Promote in Posts

CTAs appear on 99% of posts (99 out of 100), usually placed at the end after the full value is delivered. The default ask is the newsletter (93 out of 100 posts), often paired with “follow” and a repost prompt; in this dataset, posts with any CTA average about 1.23x relative engagement vs 0.58x for the single no-CTA post (directional because it’s a 1-post comparison). For higher intent, he sometimes uses keyword DMs or “repost/comment and I’ll DM the asset,” which turns engagement into both reach and a list of interested people.

[Example 1 →](#)

[Example 2 →](#)

[Example 3 →](#)

How They Promote in Profile

The profile repeats one message everywhere: leaders want ROI, and he helps them ship real, governed AI systems. The headline and banner set the promise and prompt following, the About section stacks proof (quantified outcomes + senior roles + hard industries), and the Featured section gives one clear next step: a single newsletter link—so visitors aren’t forced to choose between multiple competing offers.

[View Profile →](#)

Top 10 Posts

1



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

He argues that agentic AI is not just a chatbot connected to tools, and that many teams misunderstand what they're building. He teaches a five-layer view of the "agentic" world and focuses on the hard parts that make systems reliable in production, like guardrails, tracing, memory rules, rollback, cost control, and multi-agent coordination.

♡ 2663 💬 256 ↻ 290

2

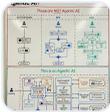


AI isn't magic.

He breaks AI into a simple layered stack from classical rules all the way to agentic systems. Each layer gets a plain-language one-line definition so non-experts can follow the progression without getting lost in buzzwords.

♡ 2506 💬 306 ↻ 318

3



Stop calling everything an AI Agent.

He challenges the habit of calling any LLM-based feature an "agent," and explains why that confusion leads to wrong expectations and weak ROI. He separates chatbots, RPA scripts, RAG-based Q&A, and true agentic systems, then defines what a real agent does (plans, acts with tools, remembers, and can coordinate).

♡ 2252 💬 297 ↻ 203

4



AI isn't magic.

He revisits the "AI is a stack" concept and explains each layer in short, skimmable sections, using simple examples like rules vs learning and "identify" vs "create." He emphasizes the shift from generative AI that responds to agentic AI that can plan and act.

♡ 2043 💬 198 ↻ 197

5



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

He explains agentic AI as more than tools attached to a model, using a five-layer model that separates foundations from the "system layer." He highlights the production concerns that make agentic systems dependable, like governance, observability, memory policies, rollback plans, cost controls, and coordination across multiple agents.

♡ 1803 💬 294 ↻ 162

6



Harvard released a full course on AI & prompting online.

He shares a free Harvard course and breaks it into a clear list of sessions, each with a quick description of what you'll learn. He also adds bonus resources to help people go deeper without searching.

♡ 1724 💬 278 ↻ 228

7



This list saves you 100+ hours of research:

He curates popular AI learning repositories and presents them as a short ranked list with star counts and quick explanations. He also shares how he filtered the list and includes a script so others can replicate the approach.

♡ 1761 💬 197 ↻ 217

8



AI isn't magic.

He argues that results depend on architecture, not just the model, and compares three system types: fixed pipelines, goal-driven agentic loops, and agentic systems grounded in company data with memory. He calls out a key weakness—ungrounded systems hallucinate—and points to what strong systems combine.

♡ 1553 💬 279 ↻ 145

9



AI's real bottleneck isn't GPUs.

He claims the limiting factor for AI isn't hardware supply, but the economics of scaling compute. He walks through a chain of large numbers (data center cost, power scale, total capex, interest rates) and ends with a strategic conclusion about unit economics and using smaller models where they add value.

♡ 1476 💬 335 ↻ 156

10



Stop calling everything an AI Agent.

He repeats the warning about "agent-washing" and clarifies the difference between chatbots, RPA scripts, RAG Q&A systems, and true agentic systems with an orchestrator. He defines agents by behaviors—planning, acting with tools, remembering, and coordinating with sub-agents—and ties it back to ROI and expectations.

♡ 1613 💬 230 ↻ 113