

TECH & AI CASE STUDY - FASHION E-COMMERCE - 2025

Visual Search That Lifted Fashion Conversion 22%

Based on a delivery engagement with Marlowe & Lark

+22%

VISUAL-SEARCH CONVERSION

+31%

REVENUE PER SESSION

-57%

ZERO-RESULT SEARCHES

9% of sessions

MOBILE ADOPTION

A fillable practice edition: the full story, 8 actions,
6 reflections and a 90-day plan - work it with your team.

Free - no email required
emersontechnology.com

MAKE THIS PRACTICE BOOK YOURS

Work it, do not just read it

This practice book turns the Marlowe & Lark case study into a working session for your team. Read the full story first - what happened, what was built and what it changed. Then tick the 8 actions you can commit to, write short answers to the 6 reflections, and finish with a 90-day plan. Every field in this PDF is fillable: open it in Adobe Acrobat Reader, Chrome, Edge or macOS Preview, type directly into the boxes and save.

YOUR NAME	ROLE
COMPANY	TEAM / DEPARTMENT
DATE	EMAIL (OPTIONAL)

MY GOAL FOR THIS PRACTICE SESSION

- 01 The story** - What happened at Marlowe & Lark - before, during and after the project.
- 02 The results** - The numbers that changed, plus industry benchmarks from public research.
- 03 The 8-action checklist** - Tick boxes as your organization completes each step. Owner and date fields included.
- 04 The 6 reflections** - Fillable questions that pressure-test your own roadmap.
- 05 My 90-day plan** - Three fillable phase cards that turn this case study into your roadmap.

RUN IT AS A 90-MINUTE TEAM SESSION

25 min READ Read the story together. Ask everyone to note the one number that surprised them most.	40 min PRACTICE Work the 8-action checklist and the 6 reflections. Assign an owner and a date to every action you commit to.	25 min COMMIT Fill in the 90-day plan. One goal per phase, one number that tells you it worked. Book a review date before you leave the room.
--	--	---

THE STORY

What happened, and why it worked

01 The Company

Marlowe & Lark is a direct-to-consumer fashion brand generating \$85M in annual revenue, selling dresses, knitwear, outerwear and accessories through its own storefront and a handful of marketplace channels. Its customer behaves in a way fashion marketers describe but rarely design for: she discovers a look on a social feed, saves it, and arrives at the brand's site wanting that outfit, not a keyword. The catalog is deep — thousands of active SKUs multiplied across seasonal colorways — and merchandise is photographed twice over, once as clean product imagery and once as lifestyle editorial, which turns out to matter enormously when machines need to learn what the brand's clothes actually look like on real bodies in real light.

By 2025, the brand's leadership could see the gap between how shoppers searched and what the site could understand. Traffic was overwhelmingly visual in origin, arriving from platforms where discovery is a photograph, and 71% of sessions were mobile, where typing long descriptive queries is at its worst. Competitors had begun shipping camera-based search, and the merchandising team had started to lose the inspiration-driven shopper — the highest-intent visitor in fashion — before the shopper ever typed a word.

02 Where Things Stood

The storefront's search ran on keyword matching against product titles and tags, and it failed precisely where fashion demand lives. Shoppers searched for looks rather than SKUs — a rust floral midi like the one on TikTok — and the keyword engine returned nothing. Eleven percent of all searches ended in zero results, and search sessions converted at 1.9% against a 3.4% site average, meaning the search bar was the site's worst-performing front door. The brand was effectively charging its most motivated visitors an extra toll: they arrived knowing what they wanted, and the site made them describe it in words the catalog did not use.

The team had tried stopgaps. Pinterest boards curated by the styling team and a styling quiz buried two levels deep in the navigation were meant to catch visual demand, and neither moved conversion nor shopper confidence in any measurable way. The quiz collected answers nobody acted on; the Pinterest boards sent traffic off-site to a platform that would happily sell the look through someone else's links. Meanwhile the zero-result list grew nightly, a running inventory of demand the catalog could see and could not serve.

Merchandising put a number on the gap: an estimated \$2.2M in annual revenue was going to competitors with camera search — shoppers who photographed a look, found it instantly elsewhere, and never returned. Deloitte Digital research has found that digital interactions influence roughly 64 cents of every dollar spent in US retail stores, a reminder that discovery experiences shape purchases well before checkout. Marlowe & Lark's discovery experience was answering a question its shoppers had stopped asking.

03 The Challenge

The challenge had two halves that had to be solved together. The first was comprehension: a model that could look at a shopper's photo of a real outfit — filtered, cropped, worn in sunlight or under bar lighting — and match it to the brand's catalog, including the colorway naming conventions merchandisers actually use. The second was retrieval design: pure image matching returns lookalikes shoppers cannot buy, because the exact dress may be out of stock, last season, or photographed in a color the model ranks first for the wrong reason. Gartner has predicted that early-adopter brands redesigning digital commerce around visual and voice search could increase digital commerce revenue by up to 30%, but that lift only materializes when the experience ends in purchasable results rather than attractive dead ends.

There was also an adoption problem hiding inside the technology problem. Pinterest reported more than 600 million visual searches per month on its Lens feature, evidence that camera-based discovery had reached mainstream consumer scale — but consumers had learned the gesture on platforms designed for it, and would abandon a camera control that hesitated, lagged or buried itself in menus. Whatever was built had to feel instant on a phone, live inside the search bar where shoppers already were, and never strand a session on a zero-result screen.

04 Discovery and Diagnosis

Emerson Technology began with a two-week discovery sprint centered on the one asset the stopgaps had ignored: the zero-result search log. The team exported months of failed queries and clustered them by intent, and the clusters told the story plainly — color-plus-garment descriptions, style references and occasion phrases dominated, and nearly every cluster described something a camera could answer faster than a text box. The audit then inventoried the visual assets available for training: studio product photography, lifestyle editorial shoots, and an underused archive of returned-item annotations that recorded why shoppers sent pieces back. Together these covered a large share of the visual variety shoppers would actually photograph.

The diagnosis also quantified the trust problem inside the funnel. Search-to-purchase conversion ran at 1.9% against the 3.4% site average, and the gap concentrated in exactly the color-and-garment query clusters the zero-result log exposed. Interviews with the merchandising team surfaced a second insight: they did not want another black-box ranking — they wanted a dashboard showing which visual intents failed, because that dashboard doubles as a demand signal. That request shaped the product: analytics for merchandisers would be a first-class feature from day one, not a reporting afterthought.

05 The Solution

Emerson Technology designed and built a visual search service on AWS, integrated into the existing Next.js storefront, that turns a shopper's photo into ranked, purchasable results in under 900 milliseconds. The architecture has three layers.

Learning the catalog's visual language

At the core is a CLIP-style dual-encoder fine-tuned in PyTorch on Marlowe & Lark's own imagery: product photography, lifestyle shots and returned-item annotations. Fine-tuning on lifestyle imagery was the decisive choice, because shoppers photograph real outfits in real light, and a model trained only on studio white-background images misreads them; training on the brand's editorial and UGC-style assets closed most of the gap between lab accuracy and field performance. The model also learned how merchandising actually names things — rust, not burnt orange; midi, not mid-length — so image results align with the catalog's language. Catalog embeddings refresh nightly into a pgvector index on Postgres, keeping the infrastructure simple at the brand's scale while ensuring results never lag the merchandising calendar.

Blending image and text retrieval

A Python FastAPI service blends image similarity with text signals — stock status, size availability and price filters — and returns ranked, purchasable results rather than raw lookalikes. When the exact photographed item is sold out, retrieval surfaces the closest in-stock alternative in the same silhouette and palette, converting what used to be a dead end into a recommendation. Fallback logic routes any low-confidence match to hybrid text-plus-image retrieval, so a shopper never hits a zero-result screen. This blend is what protected trust: shoppers forgive a near-miss they can buy; they do not forgive a beautiful result with no size in stock.

The storefront experience

The Next.js storefront gained a camera control inside the search bar itself, with example screenshots shown on first use, because placement drives adoption more than model quality in early rollouts. An experimentation layer allocates traffic for A/B tests against text search, and analytics track zero-result rates by category — the merchandiser-facing dashboard requested during discovery. Merchandisers began using the zero-result view as an early demand signal, spotting rising interest in a silhouette before the buying team had finished its own analysis.

06 The First 90 Days

Weeks one and two were spent mining and clustering the zero-result log, then selecting the launch category: dresses, the brand's strongest visual-intent category, where success would be easy to judge and embedding quality easiest to control. Weeks three through six built the training set and fine-tuned the dual-encoder, with the merchandising team grading sample matches weekly — a human-in-the-loop review that caught naming mismatches the metrics missed. Weeks seven through ten launched the pilot behind an experiment: camera search on dresses for a holdout-controlled slice of mobile traffic, measured on conversion, revenue per session and add-to-cart against text search.

The pilot readouts built the case for expansion the honest way. Visual-search sessions converted well above baseline, but the team also watched the failure modes — stale embeddings briefly surfacing discontinued colorways until the nightly refresh was enforced, and a latency regression on throttled mobile networks that visibly depressed usage until the p95 budget was restored. Weeks eleven and twelve expanded visual search beyond dresses, moved the camera control onto every category page, and wired the zero-result dashboard into the merchandising team's Monday review. Every ranking change shipped behind an experiment, so each point of lift stayed attributable to a specific decision rather than to general momentum.

07 The Results

Within two quarters, visual search went from pilot to a measurable revenue channel. Visual search now handles 9% of mobile sessions, and conversion on those sessions runs 22% above baseline search — closing, and on some clusters reversing, the historic search-to-purchase gap. Revenue per visual-search session is up 31%, add-to-cart rate on visual results climbed 27%, and zero-result searches fell 57% after image and text retrieval were unified. Search, once the site's biggest leak, has become a discovery channel the brand now markets directly.

METRIC	BEFORE	AFTER
Visual-search conversion	-	+22%
Revenue per session	-	+31%
Zero-result searches	11% of searches	-57%
Mobile adoption	-	9% of sessions

The qualitative changes compounded the quantitative ones. Merchandising uses the zero-result dashboard to spot demand before buyers do, which has quietly changed how the team talks about the search bar — from a utility that occasionally fails to a sensor on shopper intent. Bounce behavior improved as relevance improved, and the styling quiz and Pinterest stopgaps were retired without ceremony. The \$2.2M gap merchandising once estimated is no longer a leakage line in a planning deck; it is the market the brand now competes for with a feature its competitors' demos do not match.

08 What Your Business Should Do

The pattern behind Marlowe & Lark's result is repeatable, and it starts with data you already own. Mine zero-result searches first: export the queries that return no products and cluster them by intent — this free dataset shows exactly where description-based search fails and what visual search must cover first. Pilot one strong-visual category, such as footwear or dresses, because narrow scope keeps embedding quality high and makes result quality easy to judge. Train on lifestyle imagery, not studio shots alone, since shoppers photograph real outfits in real light and fine-tuning on that reality closes most of the gap between lab accuracy and field performance. Blend visual and text retrieval from the start — image similarity combined with size, price and stock signals — because pure image matching returns lookalikes shoppers cannot buy, which erodes trust quickly and raises support tickets. Design for the camera moment by putting the control inside search and showing example screenshots on first use, and A/B test against text search with a holdout group so visual search earns catalog-wide scale through measured lift rather than through demos. Refresh embeddings nightly so results never lag the merchandising calendar — stale embeddings surface discontinued items and quietly poison the experience for repeat shoppers — and watch latency like a feature, keeping end-to-end retrieval under a second on mobile networks with alerts on p95, because every extra 500 milliseconds of delay visibly depresses adoption and conversion.

Ask the questions Marlowe & Lark's team now reviews quarterly: what share of your site searches return zero results, and what do shoppers do next; how do customers describe products they cannot name, and could a camera answer faster than a text box; what fraction of your traffic arrives from visual platforms expecting a matching on-site experience; which categories would justify a narrow pilot first, and why those; do you know your search-to-purchase conversion gap between text search and the site average; and if shoppers could photograph any look tomorrow, how many catalog items could you confidently match. Work through the eight actions in this practice book with your team, and revisit those reflection questions each quarter to keep the experience honest.

REAL-WORLD PROOF & BENCHMARKS

Gartner has predicted that early-adopter brands redesigning digital commerce around visual and voice search could increase their digital commerce revenue by up to 30%. - Gartner

Pinterest reported more than 600 million visual searches per month on its Lens feature, evidence that camera-based discovery had reached mainstream consumer scale. - TechCrunch, 2019

Deloitte Digital research found that digital interactions influence roughly 64 cents of every dollar spent in US retail stores, underlining how discovery experiences shape purchases. - Deloitte Digital, 2015

THE PRACTICE PAGES

The 8-action checklist

Work through the list with your team. Tick each box as your organization completes the step, and fill in an owner and a target date so every action has a name and a deadline.

☐

1. Mine zero-result searches first

Export the queries that return no products and cluster them by intent. This free dataset shows exactly where description-based search fails and what visual search must cover first.

OWNER

TARGET DATE

STATUS

--	--	--

☐

2. Pilot one strong-visual category

Launch on a category with high visual intent, such as footwear or dresses. Narrow scope keeps embedding quality high and makes result quality easy to judge.

OWNER

TARGET DATE

STATUS

--	--	--

☐

3. Train on lifestyle imagery

Shoppers photograph real outfits in real light. Fine-tuning encoders on lifestyle and user-generated images closes most of the gap between lab accuracy and field performance.

OWNER

TARGET DATE

STATUS

--	--	--

☐

4. Blend visual and text retrieval

Combine image similarity with text signals such as size, price and stock. Pure image matching returns lookalikes shoppers cannot buy, which erodes trust quickly and raises support tickets.

OWNER

TARGET DATE

STATUS

--	--	--

☐

5. Design for the camera moment

Put the camera control inside search rather than buried in menus, and show example screenshots on first use. Placement drives adoption more than model quality in early rollouts.

OWNER

TARGET DATE

STATUS

--	--	--

☐

6. A/B test against text search

Hold out a control group and compare conversion, revenue per session and zero-result rates. Visual search earns catalog-wide scale through measured lift, not through demos.

OWNER

TARGET DATE

STATUS

☐

7. Refresh embeddings nightly

Re-encode new products and colorways every night so results never lag the merchandising calendar. Stale embeddings surface discontinued items and quietly poison the experience for repeat shoppers.

OWNER

TARGET DATE

STATUS

☐

8. Watch latency like a feature

Keep end-to-end retrieval under a second on mobile networks and alert on p95. Every extra 500 milliseconds of delay visibly depresses adoption and conversion across sessions.

OWNER

TARGET DATE

STATUS

AFTER THE CHECKLIST: OUR TOP 3 PRIORITIES

From the eight actions above, write down the three you will start first - and the single number you expect each one to move. Three priorities with owners beat eight without any.

PRIORITY 1 NUMBER IT SHOULD MOVE	PRIORITY 2 NUMBER IT SHOULD MOVE	PRIORITY 3 NUMBER IT SHOULD MOVE
--	--	--

THE PRACTICE PAGES

Reflections for your team

Use these questions in your next planning session and type short answers into the boxes. The gaps you find are usually the fastest wins.

Q1. What share of your site searches return zero results, and what do shoppers do next?

Q2. How do customers describe products they cannot name, and could a camera answer faster than a text box?

Q3. What fraction of your traffic arrives from visual platforms expecting a matching on-site experience?

Q4. Which categories would justify a narrow visual-search pilot first, and why those specific ones?

Q5. Do you know your search-to-purchase conversion gap between text search and the site average?

Q6. If shoppers could photograph any look tomorrow, how many catalog items could you confidently match?

THE PRACTICE PAGES

My 90-day plan

Finish the session by committing to a plan. Three phases, one goal each, one number that tells you it worked. Bring this page to your next leadership review - or to a free consultation with Emerson Technology, and we will pressure-test it with you.

Days 1-30 Focus: audit and alignment

PRIMARY GOAL

KEY ACTIONS

HOW WE MEASURE SUCCESS

OWNER

Days 31-60 Focus: build and prove

PRIMARY GOAL

KEY ACTIONS

HOW WE MEASURE SUCCESS

OWNER

Days 61-90 Focus: scale and measure

PRIMARY GOAL

KEY ACTIONS

HOW WE MEASURE SUCCESS

OWNER

Tip: save a copy per team so each group keeps its own answers. Revisit the reflections every quarter - the questions stay the same even as the numbers move.

BOOK THE REVIEW BEFORE YOU CLOSE THE BOOK

Plans without a review date drift. Before you close this practice book, agree on when you will come back to these pages, who will run that session, and when you will check the numbers again. If you want an outside pair of eyes, bring this page to a free consultation with Emerson Technology and we will pressure-test the plan with you.

REVIEW DATE	REVIEWED BY	NEXT CHECK-IN
------------------------------------	------------------------------------	--------------------------------------

Ready for a transformation like Marlowe & Lark's?

Emerson Technology designs and builds the data, AI, cloud and automation systems behind results like these. On a free consultation call we will map this exact playbook to your business - your data, your constraints, your numbers.

Book a free consultation

support@emersontechnology.com

or start online: emersontechnology.com/contact