

The Next Stage of AI in Loyalty

Mechanical, Thinking, or Feeling?


TBH Loyalty Exchange Live
Amsterdam • 15 October 2025

Ryan Baltrip, PhD
Exec. Director, Loyalty Science Lab
Asst. Professor, Marketing, Old Dominion
University



AI Progression: From Automation to Emotion

("The Human Touch")



Shift: from *efficiency*, to *personalisation*, toward *relationship*


Artificial Empathy

Empathy: the capacity to sense, understand, and share the thoughts and feelings of another person

Artificial Empathy: the codification of human empathy through computational models in AI agents.



Artificial Empathy




Cognitive / Thinking



Affective / Feeling

Why Artificial Empathy?


Because there's an empathy gap: 20% -> ~3%




How:

- Better **recognition**
- Credible **response**
- Fast **make-good**
- Clear **handoff**

Artificial Empathy




Cognitive Empathy (thinking, perspective-taking)	Affective Empathy (feeling, emotion-matching)
Detects customer's perspective	Recognizes positive & negative emotions
Understand needs, goals, and context	Creates expressions of care, concern, solidarity
Adjusts information & tone	Uses style & personality to sooth and encourage



**Today, AI = Strong cognitive
Simulated affective**


The Disclosure Paradox (Warmth vs Wording)

- “AI-generated” label **depresses** warmth & credibility”
- “AI-assisted” + clear benefit (e.g., faster fix) **tends to performs better**
- Label choice is **a design decision**, not an afterthought.



EU Guardrails & Where to Start

- Begin with **high-emotion, moderate-complexity touchpoints** (e.g., complaint acknowledgements, loss of benefits, delays, out of stock)
- Design with **transparency** and aim for **authenticity**
- Design for **accuracy, choice, escalation, and measurement**



When is “Feeling” AI fit for building customer trust?

If AI can only **simulate** empathy, what **conditions** would make it ***acceptable*** – and ***effective*** – for your customers in Europe?



How should it be disclosed?



What tone & actions restore loyalty?




What signals are acceptable?



When should humans take over?

Live Experiment, Part 1: EU “Feeling AI” Stress Tests

- A. Disclosure framing (authenticity)
- B. Tone vs Action (what restores loyalty?)
- C. Modality acceptability
- D. Human handoff thresholds



Experiment A:

Disclosure Framing for Building Authenticity & Trust

- Goal: identify **disclosure label** EU audiences deem most ***authentic*** to customers and label brand's find most ***acceptable*** to deploy.
- Context of the experiment:
 - Missing 12,000 points → apology + fix + prevention.
- Compare labels:
 - **AI-assisted** vs **AI-generated** vs **Human-supervised AI**.

Common Apology Text

"I'm sorry your 12,000 points didn't post after your trip. That's frustrating—especially if you planned a redemption. I've restored the 12,000 points and flagged the delay to prevent repeats."

Ai-assisted

Ai-generated

Human-supervised AI

Poll A1

Which label would your customers perceive as most authentic?

"I'm sorry your 12,000 points didn't post after your trip. That's frustrating—especially if you planned a redemption. I've restored the 12,000 points and flagged the delay to prevent repeats."

Ai-assisted

Ai-generated

Human-supervised AI

None of these

Not sure

Poll A2

Which label would your brand allow in production today?

"I'm sorry your 12,000 points didn't post after your trip. That's frustrating—especially if you planned a redemption. I've restored the 12,000 points and flagged the delay to prevent repeats."

Ai-assisted

Ai-generated

Human-supervised AI

None of these

Not sure

Experiment B:

Tone vs. Action – What Restores Loyalty?

- Goal:
 - Does **affective tone** matter less than **concrete recovery** when AI is the messenger, and how far will EU brands go with auto remedies?
- Two examples:
 - Warm apology (cognitive empathy)
 - Corrective action (less empathy; concrete restitution)

Warm Apology

"I can see how missing points after a long trip feels unfair—especially when you planned to redeem them this weekend. I'm prioritising your case and will update you within 24 hours."

Corrective Action

"Sorry for the posting delay. I've restored 12,000 points now and adjusted your tier expiry by 30 days to make this right."

Poll B1

Which response most restores trust for your customers?

Warm Apology

"I can see how missing points after a long trip feels unfair—especially when you planned to redeem them this weekend. I'm prioritising your case and will update you within 24 hours."

Corrective Action

"Sorry for the posting delay. I've restored 12,000 points now and adjusted your tier expiry by 30 days to make this right."

Warm Apology

Corrective Action

Depends on context

Neither

Not sure

Poll B2

How comfortable are you with AI auto granting compensation within pre-set rules?

1 – Very Uncomfortable ... 5 – Very Comfortable

Experiment C:

Modality Acceptability Ladder

- In EU, consent is non-negotiable.
- Goal: Identify **what signals** audiences consider acceptable for empathic AI in loyalty and where opt in is required.
- 2 polls related to 4 signals

What signals can AI read?

- **Text** sentiment (words you type)
- **Voice** tone (pitch/pace on calls)
- **Facial** camera analysis (webcam/video)
- **Behavioural** inference (clicks, purchase patterns)



Poll C1

Which signals are acceptable for AI to use in your loyalty programme?

Text sentiment

Voice tone

Facial analysis

Behaviour interface

None

Not sure

Poll C2

Which of these would require explicit opt in from customers?

Text sentiment

Voice tone

Facial analysis

Behaviour interface

None

Not sure

Experiment D:

Human Handoff Threshold

- Goal:
 - Identify **when empathic AI should escalate to a human**, and how much volume should auto escalate.
- 2 polls, 1 key question: **When must a human take over?**

Poll D1

Escalate to human when...

- Monetary loss > €50
- VIP tier affected
- Emotional distress detected
- Potential PR or regulatory risk
- Always offer human
- Other


Poll D2

What % of cases should auto escalate based on your policy?

- 0-10%
- 11-25%
- 26-50%
- 51-75%
- 76-100%

Live Experiment, Part 2:

Define **where** empathic AI is appropriate vs human-only in EU context?



Red Light / Green Light

How it works:

- Given 10 different contexts
- In pairs, pick **3 GREEN** contexts and **3 RED** contexts.
- One person from each pair, submits 3 GREENS and 3 REDS
- We'll share & discuss


10 Contexts for Red Light / Green Light

- Routine refund
- Points-expiry appeal
- Flight delay/cancellation
- Out of stock/delivery delay
- VIP downgrade

- Suspected fraud flag
- Bereavement-related cancellation
- Child/minor account
- Health/insurance query
- Legal change


What We Learned Today

- Disclosure label preference (customer vs brand).
- Tone vs action: what repaired trust.
- Signals acceptable & consent expectations.
- Human handoff thresholds.
- Red/Green context lines



Three Research-based Recommendations

- **Augment**, don't replace: *empathic triage* + *fast human escalation* for high-stakes cases.
- Start with **one high-emotion, moderate-complexity touchpoint**.
- Measure *emotions* alongside *transactions*



Thank You!

Ryan Baltrip, Ph.D., CLMP

Exec. Director, Loyalty Science Lab

Asst. Prof. of Marketing, ODU

 www.odu.edu/loyalty-science-lab

 loyaltysciencelab@odu.edu

 +1-757-683-6551



Subscribe