beefree

MANIFESTO

The Beefree and AI Manifesto

Expanding the definition of *good* digital content.

A living document · First published on June 24, 2026

Beefree · Since 2014

01

THE BEEFREE MISSION

Why we exist

Digital content is key to the way people communicate. Since 2014, Beefree's mission has been to **help more people create *good* digital content**, with a particular focus on email.

When we started, by *good* we mostly meant consistent with design best practices, responsive on mobile devices, and displaying as expected.

Over the years, we expanded the definition of *good* to include ways to comply with brand guidelines, fix problems, and produce more accessible content.

We added new dimensions to *good* with *Really Good Emails*: celebrating excellence that inspires others through a curated email gallery, research, events, and community.

We expanded the definition of *good*.

02

AI IMPACT ON OUR MISSION

Why use AI

Artificial intelligence can expand what our tools are able to do, and with it, what *good* means.

A few examples to illustrate the point. AI allows natural language interactions, no matter which language our customers speak. It can help diagnose and fix issues based on rules they set. It can show inspiration related to the content they are working on. It can flag off-brand design elements. There are many others.

We believe that, as a result, **AI can help us advance the mission** we've pursued since we started, enabling more people to create better content.

AI can help us advance our mission.

03

THE ROLE OF AI

How to use AI

We are aligned with **Human-Centered AI (HCAI)**, the discipline intent on creating AI systems that amplify rather than displace human abilities. It's the application to artificial intelligence of human-centered design, which prioritizes the needs, values, and experiences of people.

HCAI is a broad discipline. We have intentionally decided to focus on three key principles, expressed here in our own words, which can be our *constitution*.

Three principles

The Human-Centered AI principles that guide every decision we make.

1

Collaboration between humans and AI

AI in Beefree is meant to augment — not replace — people. It enhances human capabilities, handling speed and scale, while humans retain control over quality, judgment, and creative direction.

2

Explainability and traceability

AI in Beefree is explainable, allowing people to understand how it draws conclusions. It is also traceable, transparently and reliably logging its actions, which earns people's trust and may also be a legal requirement.

3

Human accountability

AI in Beefree includes meaningful human interaction, and forces human supervision when required, by law or otherwise. Logging "human-in-the-loop" events ensures accountability.

04

THE FUTURE ROADMAP

How this shapes our future

We see Beefree becoming a human-centered, AI-augmented content creation platform. Our goal is to help create better content at scale while keeping people in control. As AI expands what our tools can do, it also broadens what *good* digital content can mean.

Earlier in this document we mentioned a few practical ways in which *good* becomes broader. We are committed to exploring and bringing to market product improvements to make that possible.

OUR COMMITMENT

As we do so, we will prioritize alignment with this manifesto over *speed to market*.

The product roadmap will evolve as capabilities do, so we are intentionally not setting it here. The principles in this manifesto are our current best thinking, but we expect them to evolve too — as we learn, as the technology advances, and as the regulatory landscape develops. We commit to revisiting this document regularly and openly, and, when it changes, to communicating *how* and *why*.

APPENDIX

Glossary

Shared definitions for the key terms used throughout this manifesto.

Good digital content

NOUN PHRASE · BEEFREE TERM

Good digital content is content that does its job for everyone who receives it. Since 2014, when Beefree was founded, the definition has grown: good content renders as expected on every device and email client, follows design best practices, stays on brand, and is accessible to people with disabilities.

AI can help widen it again, as content becomes easier to translate, personalize, trace, check, and stand behind.

Human-centered AI

NOUN · HCAI · DISCIPLINE

Human-centered AI applies the principles of human-centered design to the development and deployment of AI systems. Where traditional AI research has focused primarily on capability (making systems more accurate, faster, more powerful) HCAI asks a prior question: does this system actually serve the people using it?

It was formalized as a discipline by researchers including Ben Shneiderman at the University of Maryland, whose 2022 book *Human-Centered AI* laid out the framework. HCAI calls for AI systems that amplify human ability, support human oversight, and remain controllable even as they grow more capable.

Applied to a content creation platform like Beefree, HCAI means AI that speeds up drafting, checking, and adapting campaigns, while the people responsible for the brand and the audience decide what gets sent.

More on HCAI? See Stanford HAI and the seminal Shneiderman, Human-Centered AI, OUP, 2022.

APPENDIX

Glossary

Human-centered design

NOUN · PRACTICE

Human-centered design is a design philosophy and practice that starts with the people who will use a product, not the technology behind it. It prioritizes understanding users' needs, mental models, and contexts before defining solutions.

The process is iterative: prototypes are tested with real users, insights feed back into design, and the cycle continues. It contrasts with technology-centered design, which starts with what's technically possible and works outward.

For email creation tools, that means designing around how email teams actually work: the handoffs between copywriter, designer, and approver, the brand rules they answer to, and the inboxes their subscribers read in.

More on HCD? See IDEO, What is Human-Centered Design? or the foundational ISO 9241-210 standard.

Augmentation

NOUN · APPROACH

In the context of AI, augmentation refers to using artificial intelligence to enhance what a person can do rather than to replace what they do. An augmenting system handles the parts of a task where speed, scale, or pattern recognition give it an advantage, freeing the human to focus on judgment, creativity, and context.

In email, AI can generate email or landing page copy, image alt text, and images from text descriptions; it can translate copy into other languages; it can generate ideas or refine content; it can help create and edit drafts right from AI platforms like Claude or ChatGPT, while the person keeps the judgment and the final call on what gets sent.

The distinction matters: automation removes the human from the loop; augmentation keeps them in it. Research shows that augmentation-style human-AI collaboration often produces stronger outcomes than either humans or AI working alone, particularly for creative and knowledge work.

More on augmentation? See Licklider, "Man-Computer Symbiosis," 1960 (the foundational text). For current data, see the AI Index Report 2026, Chapter 4.

APPENDIX

Glossary

Explainability

NOUN · PROPERTY

Explainability is the ability of an AI system to provide the reasons behind its output in terms the person using it can understand, and to list enough details to give clarity on what was actually done.

In email creation, an AI suggestion would come with its reason: text contrast that fails accessibility guidelines, a color outside the brand palette, a subject line flagged for length. People can then accept, adjust, or reject it on the merits.

More on explainability? See NIST, Four Principles of Explainable Artificial Intelligence (NISTIR 8312), 2021.

Traceability

NOUN · PROPERTY

Traceability is the ability to follow the origin and history of something through a process: knowing where it came from, what happened to it, and who made decisions about it along the way. Applied to AI, it means being able to see what a system did, with what inputs, and at which step.

In an email workflow, from first draft to approved send, traceability means that at each stage (drafting, editing, reviewing, approving) the people involved can see what the AI generated, what a human changed, and what was left untouched. That visibility is what makes meaningful judgment possible: you can only truly own a decision if you know what you're deciding on.

APPENDIX

Glossary

Human accountability

NOUN · PRINCIPLE

Human accountability means that a person remains responsible for the outcomes an AI system produces. It's the principle that AI can assist, recommend, or automate, but when something goes wrong, there is always a human who owns the result.

This matters because AI systems can fail in ways that are hard to predict, and diffuse responsibility ("the AI decided") is both ethically insufficient and, in many jurisdictions, legally untenable.

In a content creation platform like Beefree, that person is whoever approves the content: if an AI-drafted email campaign goes out with a wrong price, a broken link, or a claim the brand can't support, the approver answers for it, which is why they need to see what the AI did before it ships.

Human-in-the-loop

NOUN · PRACTICE

Human-in-the-loop describes a workflow where a person reviews, corrects, approves, or overrides AI output at defined points before it takes effect. How much review is needed depends on the stakes: a layout suggestion may need none, while an email about to reach thousands of inboxes needs a person's sign-off.

Recording when these checks happen, and who made them, is what turns oversight into accountability.

More on human oversight? See Regulation (EU) 2024/1689, Article 14, and IBM's What is human-in-the-loop?.

Expanding the definition of *good* digital content.

beefree

From spark to confident creation.

beefree.io