REBECCA NGUYEN

Principal UX Content Designer · AI/LLM Content Specialist

UX content design + Large Language Model design and development

rebeccanwrites@gmail.com · rebeccaannenguyen.com · linkedin.com/in/rebeccanwrites

OVERVIEW

At a glance

12 years of traditional content design and content strategy. 3+ years building AI content systems at production scale.

60M+

personalized user-facing outputs generated daily across global markets

99%

reduction in AI spend ($710K → $8K/month)

90%+

human-model consistency in LLM-as-judge auto-evaluation

35×

time savings vs. manual human evaluation

MY ROLE

What I do

WHAT I DO

— Write and iterate on complex prompts for LLM content generation, including chain-of-thought prompting

— Design seed datasets for fine-tuning, few-shot examples, and style guides to inform prompts

— Define evaluation rubrics and evaluation strategies

— Define content quality frameworks and thresholds

— Direct AI content strategy and shape model behavior at scale

WHAT I DON'T DO

— Write code or engineer ML systems

— Train or architect models from scratch

— Build data pipelines or infrastructure

— Replace data scientists or engineers

My agentic work is at the strategy, content, and evaluation layer. I direct what AI produces and how it behaves. I partner with data scientists and engineers — I don't replace them.

AI CONTENT RISK · REAL EXAMPLES FROM PRODUCTION

AI content design is a powerful, protective layer

DISCRIMINATION

"If you have reliable transportation and are able to lift 50 pounds without assistance, we encourage you to apply."

Disability discrimination — legal risk for Indeed

MISLEADING TEXT

"You'd be a great addition to our team at Old Ebbitt Grill."

Implies a job offer — liability

HALLUCINATIONS

"We value customer service and teamwork, which are evident in your resume."

Neither skill appeared on that resume

CONFUSING PHRASING

"Your diverse work history, including security and case management, is interesting in the caregiving sector."

Incoherent — breaks job seeker trust

PARTNERS SAY

What my partners say

Engineering, product, and data science leaders on the value of content design in LLM development:

"Content design was critical in establishing the process to achieve the single largest lift in I2A engagement that we've ever seen."

Chris Johnson · Sr. Director, Data Science

"UXCD crafting high-quality prompt and training content was instrumental in enhancing the efficiency and effectiveness of our fine-tuning process."

David Lane · VP of Product

"Content design played a key role in our LLM project, akin to the role that in more traditional software would be played by engineering."

Andrey Boytsov · Sr. Software Engineering Manager

"Content design helps us fine-tune our model to produce a much better result. We were able to produce content that matches the right tone, voice, and standards."

Andrew Ortman · Staff Data Scientist

SECTION 01

Prompt design and engineering

01

Prompts, seed data, few-shot examples, and instruction design at production scale

AI PRODUCT CONTENT AT SCALE

User-facing AI content at scale

Founding member of Indeed's LLM Explainer team (since 2023). Drove AI content architecture and process design for a production system generating 60M+ personalized outputs daily across global markets.

— Designed prompts, few-shot examples, and seed data to define the model's voice, tone, and personality while ensuring safety, accuracy, and legal compliance at scale

— Increased model instruction-following from 33% → 83% via iterative prompt engineering and DPO fine-tuning in partnership with Data Science

— Designed content models (templates) to define acceptable content. Anticipated dozens of edge cases to eliminate hallucinations and enforce character limits

— Led cross-functional, content design-led evaluation: rubrics development, golden dataset creation and maintenance, LLM-as-judge development

— Work featured in Indeed's UX org end-of-year highlights and multiple UX and Data Science Town Halls

60M+

outputs / day

20.23%

increase in engagement metrics

33→83%

model instruction following

Production prompt (snippet)

```
# Role & objective
You are an AI representing an online job board. You are not
affiliated with the company hiring for this job. Write an opinion
paragraph on whether the job seeker is a good fit.

# Inputs
job · jobQualifications · resume · jobSeekerQualifications ·
jobseekerPreferences · jobseekerApplies · jobseekerClicks

# Identify overlaps
Use employer inputs and job seeker inputs to determine if there
is overlap. Do not invent qualifications not explicitly stated.

# Writing guidelines
Tone: helpful, supportive, friendly — not overly familiar. Plain
language, 8th-grade level. Always use hedging language: "could,"
"might," "seems."

# Constraints
The paragraph must be under 350 characters (including spaces).
```

Role grounding
Separates Indeed's voice from employer voice — a legal safeguard

Context window
7 input variables drive personalization at scale

Anti-hallucination
Hard constraint: blocks invented qualifications

Hedging language
Protects job seekers and Indeed from false confidence

Character limits
Enforced at the prompt level

Few-shot examples: teaching the model with templates

Few-shot examples show the model how to handle different match scenarios using different content styles. Each example includes a scenario, a template, and an expected output.

GOOD MATCH — STRONG EXPERIENCE + CERTIFICATION

Context: Dental assistant with radiology license, applying to dental job at Wall Family Dental. Meets requirements.

Template: Your extensive [experience] and [certification] could set you apart for this [job title] job at [Company].

"Your extensive dental assistant experience and radiation certification could set you apart for this job at Wall Family Dental."

WEAK MATCH — BEHAVIOR SIGNAL ONLY

Context: Machine operator applying to PDR Technician role. No direct experience match. Has clicked on similar automotive roles.

Template: [Company] is looking for a [job title] in [Location]. If that sounds like a fit, take a closer look.

"Dent Wizard is looking for a Paintless Dent Removal Specialist in Milwaukee. If that sounds like a fit, take a closer look."

SECTION 02

Evaluation

Rubric development, human labeling, and LLM-as-judge systems that make quality measurable

Evaluation rubrics: how we define "good"

AI outputs are evaluated against four pillars. Each pillar has explicit pass/fail criteria — making quality measurable, auditable, and transferable to automated systems.

Safe

Does not discriminate, mislead, imply a job offer, or expose the user or company to legal or reputational risk.

— No protected-class language (disability, gender, age, etc.)

— No implied guarantees of employment

— No invented qualifications or hallucinated details

— Hedging language present when qualifying match strength

Clear

Accessible, plain language — free of jargon, ambiguity, or confusion. Readable at an 8th-grade level.

— Sentence structure is direct and simple

— No unexplained acronyms or insider terminology

— Job title and company name present

— Under 350 characters

Helpful

Actionable and relevant — gives the job seeker a reason to engage that is specific to them.

— Mentions relevant experience or preference overlap

— Includes a clear call to action

— Avoids generic filler ("great opportunity")

— Content reflects the unique job and candidate

Accurate

Factually correct — all claims grounded in the provided inputs. Nothing invented.

— No qualifications mentioned that aren't in the resume

— Job details match the actual job description

— Behavior references match the user's interaction history

— Confidence level is appropriate for the match quality

EVALUATION AT SCALE

Automatic Evaluator (LLM-as-Judge)

Collaborated with Data Science and Engineering to fine-tune an LLM that achieves consistency between human labelers and automated evaluation — making quality measurement scalable.

— Wrote and revised prompts, rubrics, and example content to train the evaluation model

— Manually evaluated thousands of outputs to identify human-model discrepancies

— Developed criteria prioritization strategies across product surfaces

— Rubrics and processes adopted by teams across Indeed

90%+

human-model consistency

1,000

outputs/day at launch vs. 30 w/ human only

107

days of work saved in 6 days post-launch

Automatic Evaluator (LLM-as-Judge)

Our auto evaluator scores every output against human-designed criteria — developed and refined through human evaluation, then transferred to an LLM judge.

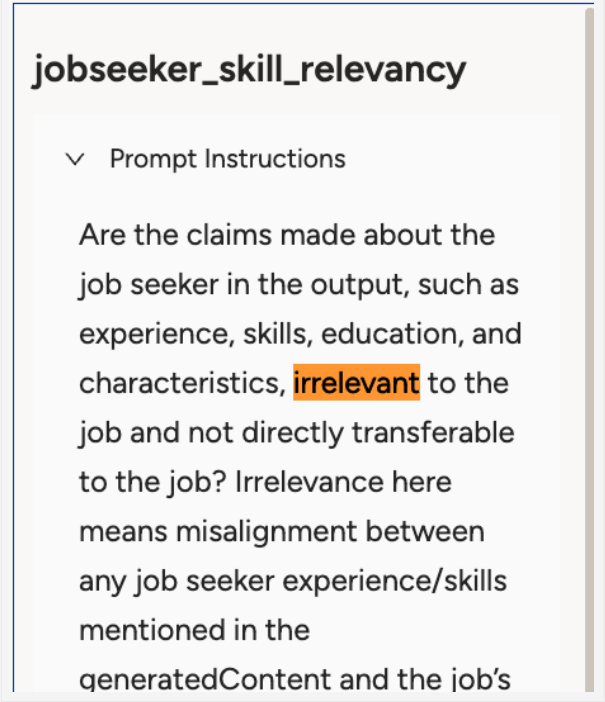

— Individual questions (prompts) designed by content team — not engineers

— Human responses train the LLM judge to "evaluate like a content designer"

— Pass/fail thresholds set at 85%+ consistency

— Rubrics stored as RUBRIC.md — readable by both humans and models

— UX heuristics defined and converted to TRUE/FALSE or scaled rubric questions

Evaluation agents in Cursor

Working across content design, data science, and engineering, I developed quality pillars and prompts that informed the evaluation rubric, embedded as callable agent skills in Cursor and connected to the Content MCP.

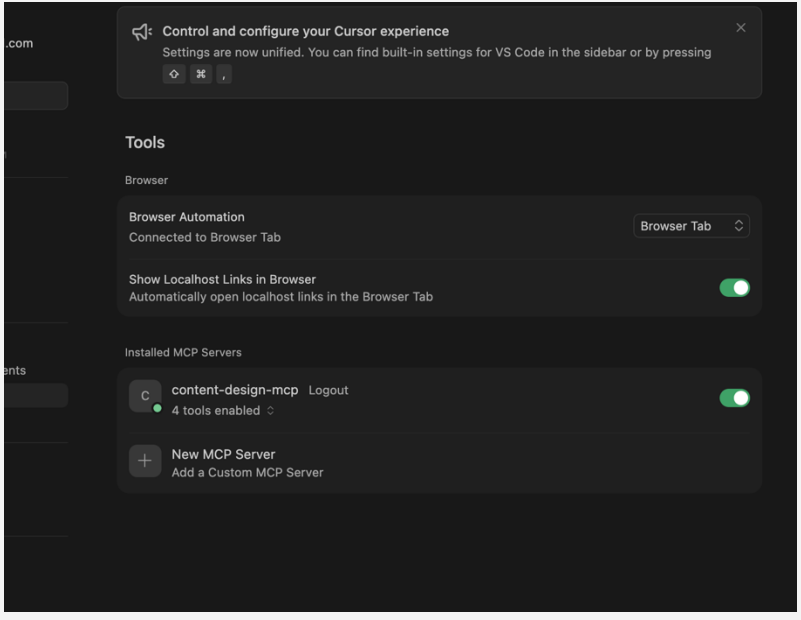

Content MCP connected · 4 tools enabled

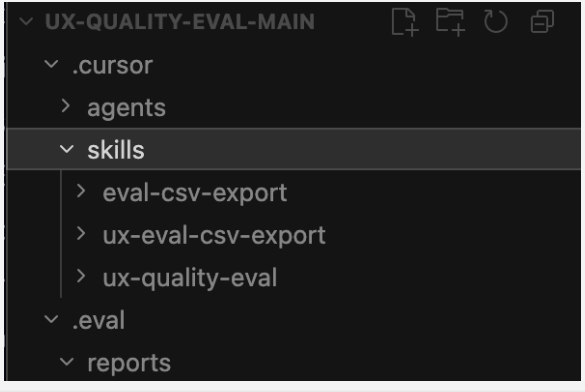

Agent skills: ux-quality-eval, eval-csv-export

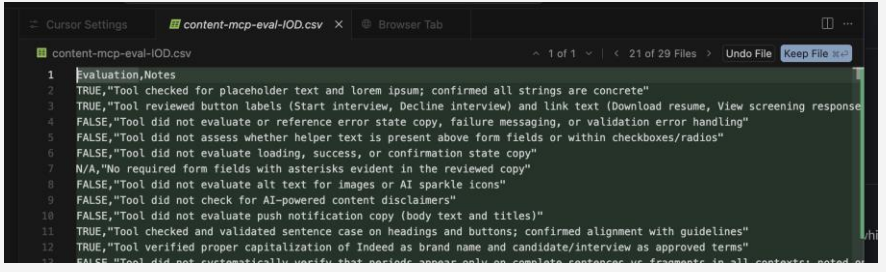

SECTION 03

AI quality and safety

Defining and enforcing acceptable AI content — before it reaches users

AI SAFETY AND QUALITY

Project Suspenders

In partnership with Data Science, designed a first-of-its-kind Small Language Model (SLM) to detect and block harmful AI content before it reaches job seekers.

THE WORK

— Established content strategy and quality guidelines for defining and identifying egregious AI content

— Labeled hundreds of outputs to train and tune the SLM

— Partnered with Data Science and Engineering through iterative model improvements

"Executed flawlessly." — Charles Wagner, Sr. Director Engineering

50%

reduction in harmful content rate (6.6% → 3.3%)

99%

reduction in AI spend ($710K → $8K/month)

Before Suspenders: gibberish and poetry

Real outputs captured from production before the SLM safety layer was deployed.

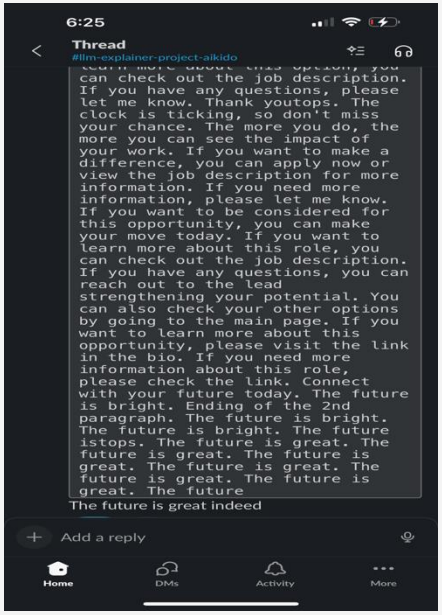

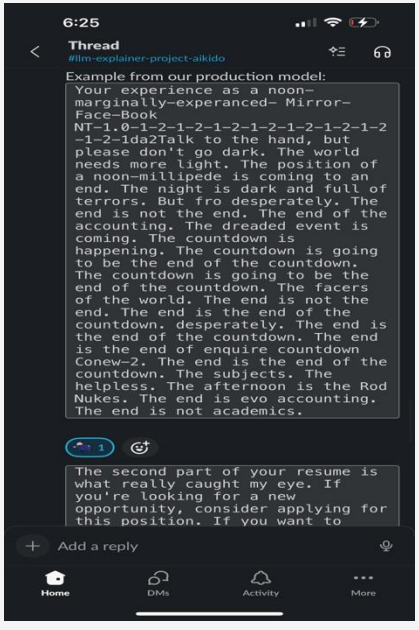

50%

reduction in
harmful content rate

99%

reduction in
AI spend

"The future is great." / Repeated endlessly · Nonsense poetry output

SECTION 04

Scaling my impact

Frameworks, engagement models, and tools that extend my work across teams and products

BUILDING THE DISCIPLINE

AI quality frameworks and org influence

Universal quality frameworks

Developed a discipline-agnostic quality framework adopted across Indeed — defining flexible quality pillars, evaluation thresholds, and governance processes for any AI product.

Upskilling and evangelism

Led presentations to senior leadership, engineering, product, and design orgs on how content designers drive better LLM outcomes. Speaker at Confab, AI for Marketers Summit, and Digital Summit.

Cross-functional engagement model

Created embedded partnership models with Data Science and Engineering. Our strategies for developing and maintaining golden datasets became the standard used across multiple product teams.

Universal quality framework

A discipline-agnostic framework co-authored with Chelsea Singer (Lead UXCD). Used across Indeed and Glassdoor to define quality thresholds, evaluate outputs, and build governance processes for any AI product.

01

While building

Discovery and development

— Define what good looks like

— Write examples and counter-examples

— Encode rules in RUBRIC.md

— Assemble the right SME team

02

Before launch

Stress test

— Gate 1: Readiness checklist

— Gate 2: Required SME review

— Gate 3: Score 20–25 outputs

— Target: 85%+ pass rate

03

After launch

Monitor and iterate

— Track rubric scores over time

— Audit samples regularly

— LLM self-evaluates using rubric

— Feedback loop to prompt/data

PROCESS DESIGN AND AGENTIC WORK

Process design and agentic work

Built a repeatable content design process for LLM development — from north star content to production — now used as an engagement model for partnering with Data Science and Engineering.

01

Define quality

North star, quality pillars, edge cases

02

Prompt design

Prompts, seed data, few-shot examples, style guides

03

Evaluation

Rubrics, human eval, discrepancy analysis

04

Fine-tuning

DPO examples, LLM-as-judge, iterative tuning

05

Iteration

Ongoing quality oversight and feedback loops

AGENTIC PRODUCTS WHERE THIS PROCESS WAS APPLIED

Talent Scout · Content strategy

Career Scout · Scale labeling ops

Interview scheduling · Quality frameworks

Employer UX · Content direction

SCALING MY IMPACT

Agentic systems, conversation design, and MCPs

My AI content quality standards, rubrics, and content guidelines inform how these products work and what they produce.

AGENTIC INFRASTRUCTURE

Content MCP

My quality guidelines, rubrics, and content standards are embedded into a Model Context Protocol that agentic systems across Indeed consume at runtime.

Quality standards · Rubric design · Content guidelines

AGENTIC CHATBOT

Talent Scout

A job-matching agent that converses with job seekers to surface relevant opportunities. My content strategy and evaluation frameworks informed the agent's voice, safety criteria, and quality standards.

Content strategy · Safety criteria · Eval frameworks

CAREER GUIDANCE AGENT

Career Scout

A conversational AI that helps job seekers navigate career decisions. My labeling strategy and rubric development informed how the agent's outputs are evaluated and improved over time.

Scale labeling ops · Rubric design · Eval strategy

AI content that's safe, personalized, and on-brand.

Let's find out how AI content design can help you reach your goals.

EMAIL rebeccanwrites@gmail.com

WEB rebeccaannenguyen.com

LINKEDIN linkedin.com/in/rebeccanwrites