

REBECCA NGUYEN

Principal UX Content Designer · AI/LLM/SLM Content Specialist

UX content design and strategy + Large Language Model design and development

reccanwrites@gmail.com · rebeccaannnguyen.com · linkedin.com/in/reccanwrites

At-a-glance

12 years of traditional content design and content strategy. 3+ years building AI content systems at production scale — from prompt and context engineering to fine-tuning and evaluation.

60M+

personalized user-facing
outputs generated
daily across global markets

99%

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

90%+

accuracy achieved in our
LLM-as-judge
auto-evaluation system

35×

time savings when
evolving from human
evaluation to automated
evaluation
(from 30 outputs/day to
1000+ outputs/day on
launch)

What I do (and what I don't)

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 closely with engineers and data scientists for the technical implementation of my strategies, but I don't replace them.

AI content design is a powerful, protective layer against:

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

What my partners say

Engineering, product, and data science leaders weigh in 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 observed (in my 5 years at Indeed) based solely on content improvements."

Chris Johnson · Sr. Director, Data Science



"UXCD crafting high-quality prompt and training content was instrumental in enhancing the efficiency and effectiveness of I2A."

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 a combination of PM, QA, and DataSci."

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 our style and quality expectations."

Andrew Ortman · Staff Data Scientist

01

Prompt engineering

Prompts, seed data, few-shot examples, and fine-tuning at production 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 product suite within a 2-sided marketplace. Defined content strategies for personalized, single-turn LLM outputs, including 60M+ daily emails across global markets.

- Designed prompts, few-shot examples, and seed data to define the model's voice, tone, and personality while ensuring safe, accurate, relevant AI content.
- Increased model instruction-following from 33% → 83% via iterative prompt engineering + DPO fine-tuning in partnership with Data Science
- Designed content models (templates) to define acceptable content. Anticipated dozens of edge cases to eliminate hallucinations. Reduced inaccurate content by 50%.
- Led cross-functional, content design-led evaluation: rubrics development, golden dataset creation and maintenance, LLM-as-judge development and refinement, and human labeling.
- 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 opening paragraph for an outreach email to a job candidate.

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 present in job seeker inputs.

Writing guidelines

Tone: helpful, supportive, friendly – not overly familiar. Plain language, 8th-grade level. Always use hedging language: "could be a good match", "might be a good fit". Avoid disqualifying statements.

Constraints

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

- Separates Indeed's voice from employer voice — a legal safeguard
- 7 input variables drive personalization at scale
- Hard constraint: blocks invented qualifications
- Protects job seekers and Indeed from false confidence or implied offers
- 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 context, a template, and an expected output.

SCENARIO: GOOD MATCH — STRONG EXPERIENCE + CERTIFICATION

- Context:** Dental assistant with radiology license, applying to dental job at Wall Family Dental. Meets requirements. Similar jobs recently considered.
- Template:** Your extensive [experience] and [certification] could set you apart for this [job title] job at [Company]. The position is similar to other jobs you recently considered on Indeed...
- Output:** "Your extensive dental assistant experience and radiation certification could set you apart for this job at Wall Family Dental. The position is similar to other jobs you recently considered on Indeed, and it meets your preference for a day shift."

SCENARIO: 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]. Ideal candidates will have [Requirements]. If that sounds like a fit, feel free to apply now.
- Output:** "Dent Wizard is looking for a Painless Dent Removal Specialist in Milwaukee. Ideal candidates will have at least one year of prior dent removal experience. If that sounds like a fit, feel free to apply now."

02

Evaluation

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

Evaluation rubrics: How we define and measure "good"

AI outputs are evaluated against four pillars. Each pillar has explicit pass/fail criteria — making quality measurable, not subjective.

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

Automatic Evaluator (LLM-as-Judge)

Collaborated with Data Science and Engineering to fine-tune an LLM that achieves consistency between human labelers and automated quality evaluation — the first of its kind at Indeed.

- 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/day w/human only

107

days of work saved
in 6 days following launch

Automatic Evaluator (LLM-as-Judge)

Our auto evaluator (LLM-as-Judge) scores every output against human-designed criteria — developed and refined through hundreds of hours of manual evaluation.

jobseeker_skill_relevancy

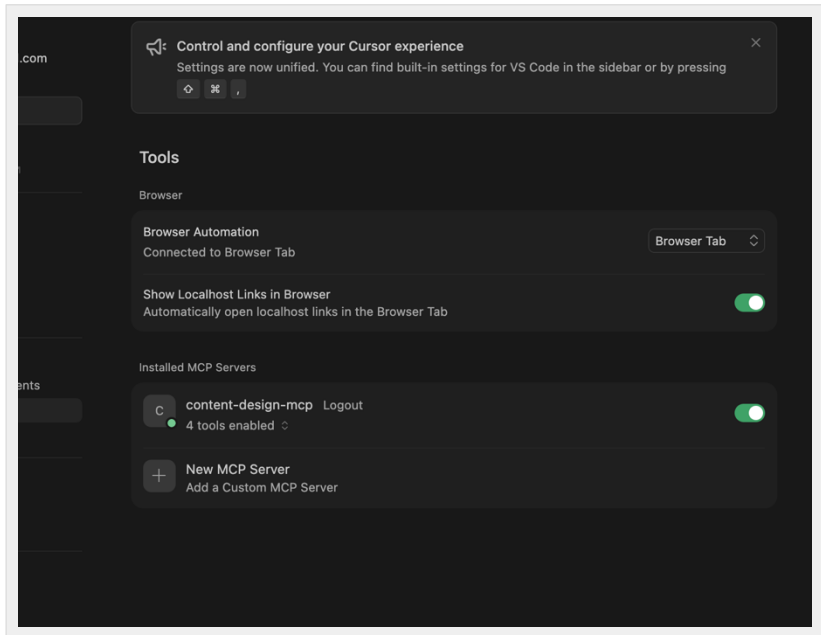
∨ Prompt Instructions

Are the claims made about the job seeker in the output, such as experience, skills, education, and characteristics, **irrelevant** to the job and not directly transferable to the job? Irrelevance here means misalignment between any job seeker experience/skills mentioned in the

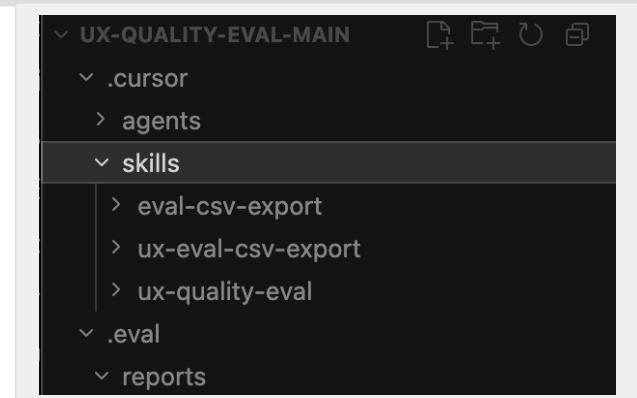
- — UX heuristics defined and converted to TRUE/FALSE or scaled rubric questions
- — Individual questions (prompts) designed by content team — not engineers
- — Human responses train the LLM judge to ‘evaluate like a content designer’
- — Each question requires its own strategy. Wording is tailored to direct model behavior and improve accurate responses. In this example, “is anything irrelevant?” out-performed “Is everything relevant?”
- — Pass/fail thresholds set at 85%+ consistency
- — Rubrics stored as RUBRIC.md — readable by both humans and models

Evaluation agents in Cursor

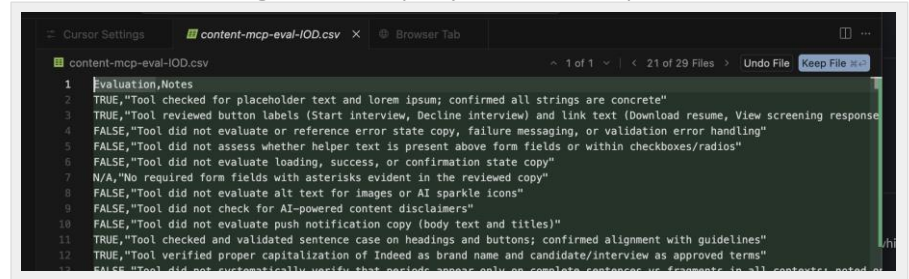
Working across content design, data science, and engineering, I developed quality pillars and prompts that informed the evaluation rubric, embedded as callable agent skill in Cursor and connected to the Content MCP. Together with another content designer, I ran evaluation prompts through the agent to assess whether it returned correct outputs, then compared those results against manual human evaluation.



Content MCP connected · 4 tools enabled



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



Evaluation output from agent invocation

03

AI quality and safety

Defining and enforcing acceptable AI content — before it reaches users

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 reached job seekers. Evaluated production content across a 60M/day global email stream.

- 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 DataSci and Eng 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)

AI quality frameworks & org influence

Universal quality frameworks

Developed a discipline-agnostic quality framework adopted across Indeed — defining flexible quality pillars, evaluation rubrics, and minimum quality thresholds for AI-powered tools and agentic systems.

Upskilling and evangelism

Led presentations to senior leadership, engineering, product, and design orgs on how content designers drive better LLM results. Developed internal guidelines, checklists, and LLM process documentation adopted across teams, embedded into our Content MCP, and consumed by agentic systems across tools.

Cross-functional engagement model

Created embedded partnership models with Data Science and Engineering. Our strategies for developing and maintaining golden datasets, rubrics, and evaluation tools are now used by teams across Indeed and Glassdoor.

Universal quality framework

A discipline-agnostic framework co-authored with Chelsea Singer (Lead UXCD). Used across Indeed and Glassdoor to define quality standards for AI-generated tools and outputs.

01

While building

Discovery & 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 & iterate

- Track rubric scores over time
- Audit samples regularly
- LLM self-evaluates using rubric
- Feedback loop to prompt/data

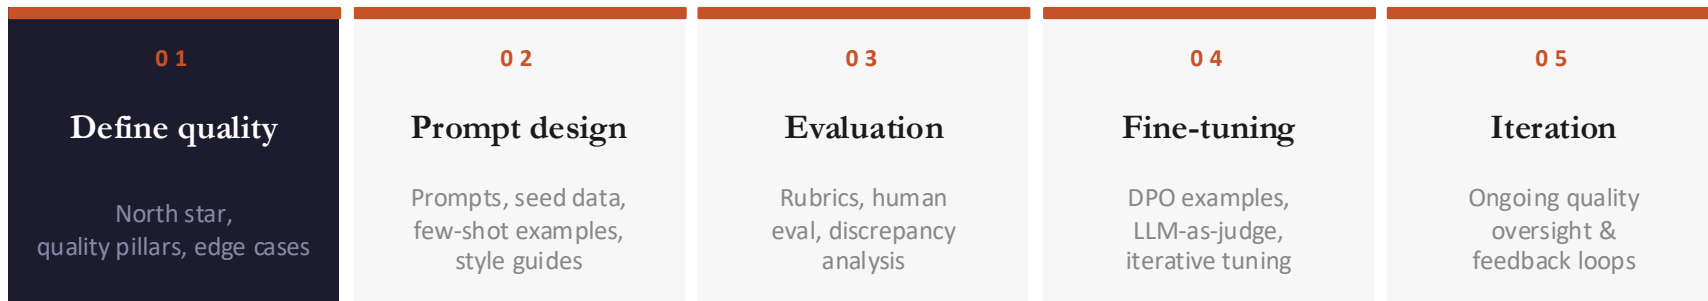
04

Scaling my impact

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

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 across multiple teams and agentic products at Indeed.



Agentic products where this process was applied

Content strategy, prompt direction and quality oversight for conversational agents



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 team's approach to 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 rebeccaannnguyen.com

LINKEDIN [linkedin.com/in/rebeccanwrites](https://www.linkedin.com/in/rebeccanwrites)