BLUF: The HireMe BlueJay Coach (HBC) is an AI interview preparation system built on Hunch.tools that help Johns Hopkins MS BAAI students prepare for job interviews. Using candidate materials and job details, it generates customized interview questions with suggested responses, guidance, recommendations, and fit assessments. Enhanced with persona adaptation, ethical guardrails, cross-cultural communication support, and automated response evaluation, it employs Claude 3.7 Sonnet (thinking), Claude 3.5, and OpenAI's o1 models for complex reasoning tasks.

Tool Link: https://app.hunch.tools/app/tool/18yxQj

System Design: HireMe BlueJay Coach (HBC) on Hunch.tools

Overview: The HireMe BlueJay Coach (HBC) is a prototype agentic AI system developed on Hunch.tools. It aims to provide personalized interview coaching for Johns Hopkins University MS BAAI students, helping them prepare for roles in fields like business analytics and AI. HBC employs a visual, node-based workflow within Hunch.tools to process user inputs and generate comprehensive interview preparation materials.

Architecture: The HBC system architecture consists of interconnected nodes within the Hunch.tools interface:

1. **Input Nodes:** Content blocks collect user-provided data for the nodes "Candidate Resume," "Cover Letter," "Job Description," and "Interviewer Persona."
2. **Processing Nodes and Data Flow**
    a. **Interview Questions & Responses (Claude 3.7 Sonnet)**: Creates 10 tailored questions, verbal responses through the STAR method, and non-verbal cues based upon content from all input nodes.
    b. **Interview Performance Review (Claude 3.7 Sonnet):** Produces strategic interview advice based on all input nodes.
    c. **Recommendations (Claude 3.7 Sonnet):** Suggests relevant learning resources and activities based upon the job description input node.
    d. **Answer Evaluator (Claude 3.7 Sonnet):** Evaluates the clarity, relevance, STAR, impact, and conciseness of the generated verbal responses of the Interview Questions & Responses Output node
    e. **Bias Check (Claude 3.5 Sonnet):** Reviews generated questions from the Interview Questions & Responses Output node for potential biases and unfair language.
    f. Cultural Context Guidance (Claude 3.5 Sonnet): Produces advice on potential cross-cultural communication dynamics within U.S.-based interviews, drawing from the Interview Persona, Candidate Resume, and Job Description input nodes.
    g. Fit Assessment (Claude 3.7 Sonnet): Evaluates candidate-role alignment based on the Cover Letter, Job Description, and Candidate Resume input nodes, providing strengths, possible gaps, and a likelihood estimate.
    h. HBC Summary (OpenAI o1): Consolidates outputs from all processing nodes into a condensed, concise final report.
3. **Output Nodes (Content Blocks):** Displays the generated content from each processing node.

Prompt Design and Ethical Guardrails

***Prompting Strategy*:** Each AI Task node contains a prompt that clearly defines the AI's role and objective, details the required tasks with specific instructions, includes explicit ethical guardrails, and defines the desired output format to ensure structured and relevant results.

Key Prompts: For conciseness, two prompts will be written in full.

1. **Interview Questions & Responses:**

***Objective*:** Act as the HireMe BlueJay Coach (HBC). Your goal is to generate interview preparation materials for a candidate applying for a specific job, based on their resume, cover letter, and job description, tailored to a specific interviewer persona, including adjusting question style and focus based on that persona.
Task: Generate exactly 10 interview questions the specified interviewer persona would likely ask the candidate for this job. Base these questions on the details found in the candidate's resume, cover letter, and the requirements in the job description.
Adapt Question Style to Persona: Explicitly consider the `Interviewer Persona` description when formulating the questions:
If the persona suggests a Hiring Manager or Director level, include a mix of strategic questions (about vision, leadership, business impact) and behavioral questions.
If the persona suggests HR, lean more towards behavioral questions focusing on cultural fit, teamwork, handling conflict, and alignment with company values.
If the persona suggests a Peer or Technical role, include more questions on specific skills, tools, processes, and day-to-day execution relevant to the job description.
Regardless of persona, ensure questions probe the experiences detailed in the resume and cover letter as they relate to the job description.
For each generated question:
a. Provide a concise, strong suggested Verbal Response for the candidate, using the STAR method (Situation, Task, Action, Result) where appropriate. Draw specific examples and context from the resume and cover letter. Keep the response focused and impactful.
b. Provide concise suggested Non-Verbal Cues for the candidate to use while delivering the answer. These cues should engage the interviewer's "limbic brain" by conveying confidence, warmth, and engagement (e.g., posture, eye contact, tone modulation, thoughtful pauses). Keep suggestions brief.
Ethical Guardrails:
Ensure all questions and suggestions are professional, fair, and unbiased. Avoid discriminatory assumptions based on any demographic information. Focus solely on the professional qualifications and experiences presented in the provided documents.
Output Format:
Present the output as a numbered list (1 to 10). For each number, include:
The Question.
The suggested Verbal Response.
The suggested Non-Verbal Cues.

2. **Bias Check Prompt:**

***Objective*:** Act as an AI Ethics Reviewer for the HireMe BlueJay Coach (HBC). Your goal is to analyze the generated interview preparation materials for potential bias, stereotypes, or unfair language, ensuring alignment with ethical AI principles for hiring assessments. Carefully review the entire text of the Interview Questions output, examining all 10 questions, verbal responses, and non-verbal cues. Analyze the text specifically for:
a. Potential Bias: Look for any language or framing that could reflect bias based on gender, race, ethnicity, age, disability, cultural background, or other protected characteristics.

b. ***Stereotypes****:* Identify any reliance on harmful or inaccurate stereotypes in the questions or suggested responses/cues.
c. ***Discriminatory Language****:* Check for any language considered discriminatory, exclusionary, or unfair to certain groups.
d. ***Unwarranted Assumptions****:* Look for questions or suggestions based on assumptions not supported by the professional context provided in the original resume/job description/cover letter.
Ethical Context*:* Fairness and equity are critical in AI-driven hiring tools. This check aims to identify potentially problematic content generated by the AI.
Output Format*:*
Provide a concise summary of the review.
* If potential issues are identified, list the specific question number(s) and the potentially problematic text, briefly explaining the concern (e.g., "Question 5 response suggests..., which could rely on [type] stereotype.").
* If no potential bias, stereotypes, or unfair language are detected after careful review, state clearly: "No potential bias or fairness issues detected in the generated content."

3. **Interview Performance Review Prompt:** Guides the AI to synthesize the resume, cover letter, job description, and persona to produce strategic interview advice. This advice includes themes to emphasize, addressing interviewer priorities, mindset and impression strategies, and adherence to company values.
4. **Recommendations Prompt:** Tasks the AI to analyze the job description to suggest relevant professional development resources across categories including readings, videos, courses, and activities which would amplify candidate's attractiveness in that role and field**.**
5. **Answer Evaluator Prompt:** The AI will review the previously generated verbal responses for all 10 interview questions based on clarity relevance, SSTAR method usage, impact, and conciseness.
6. **Cultural Context Guidance Prompt:** This prompt specifically instructs the AI to generate advice on appropriately handling cross-cultural communication differences, such as directness, non-verbal cues, and self-promotion. The prompt is relevant for international students interviewing in the U.S., ensuring that framing is respectful.
7. **Fit Assessment Prompt:** Guides the AI to objectively compare the candidates resume and cover letter against the job description. The prompt identifies strengths, potential gaps, and provides a realistic likelihood estimate of low, moderate, or strong for advancement, based only on the documents provided.
8. **HBC Summary Prompt:** Instructs the AI to consolidate the outputs from all preceding functional nodes into a single, cohesive summary report.

Ethical Guardrails: Explicit "Ethical Guardrails" sections were included in nearly all prompts, instructing the AI to avoid discriminatory assumptions, focus solely on professional qualifications presented in the documents, and maintain fairness. The "Bias Check" node provides an additional review layer. The "Cultural Context Guidance" node directly addresses fairness and inclusivity for the target MS BAAI student population. The system avoids making definitive hiring predictions, focusing instead on preparation and document-based assessment (Coley 2024).

Language Model Rationale

Model Selection: Throughout development, several LLMs were tested through the Hunch.tools interface, including adjacent LLMs developed by Anthropic, OpenAI, Google, and Grok. The final implementation uses Anthropic's Claude 3.7 and 3.5 models for all processing nodes except for the HBC Summary node, which uses OpenAI's o1 model.

Justification: Claude 3.x models (3.5/3.7) handled complex tasks, such as question generation with STAR responses and persona adaptation, evaluations, and bias checks due to consistency and contextual relevance. OpenAI's o1 was used for final summary reporting, leveraging its synthesis capabilities and context window to ensure consolidation (Goleman, 2006).

Evaluation Metrics

Proposed Metrics:

1. **Functional Completeness:** Pass/fail checklist for all required outputs
2. **Question Relevance:** 1-5 human score on alignment with inputs
3. **Response Quality:** Percentage of responses correctly using STAR method
4. **Non-Verbal Cue Actionability:** 1-5 human score on clarity/practicality
5. **Bias Audit Score:** Quantitative score from bias check results
6. **Guidance Relevance:** 1-5 human score on actionability of recommendations
7. **Fit Assessment Realism:** 1-5 human score on objectivity versus inputs

Evaluation Process: Evaluation involves running HBC with a diverse set of sample inputs, such as different resumes, job descriptions, and personas, and having human reviewers score the outputs against the defined metrics using standardized rubrics.

Conclusion

The HireMe BlueJay Coach prototype combines Anthropic Claude 3.5, Claude 3.7, and OpenAI GPT-4o models to create an AI interview preparation tool that exceeds requirements. It generates customized questions, responses, guidance, and recommendations based on candidate and job information. Added features include answer evaluation, bias checking, cultural communication advice, and fit assessment. Potential improvements could include interactive answer practice and prompt refinement based on user feedback.

References:

Coley, C. (2024). *Gender biased algorithms and replacing human interaction: A deep dive into the ethical considerations of using AI in your recruitment process*. Hydrogen. https://www.hydrogengroup.com/gender-biased-algorithms-and-replacing-human-interaction-a-deep-dive-into-the-ethical-considerations-of-using-ai-in-your-recruitment-process

Goleman, D. (2006). *Social intelligence: The new science of human relationships*. Bantam.