

Decision-Support and Wayfinding Graph

Knowledge packet for Decision-Support and Wayfinding Bot

Chapter 8: The Skeleton Layer - Design Wayfinding, Clarity, and Decision Support

Purpose

Organize choices, tradeoffs, signals, noise, confidence gaps, corrections, and next-best support so the guest can understand the decision rather than simply receive a recommendation.

Central question

What should the user understand next?

AI Employee Experience Architecture Check

Layer	What to check
Customer Reality	What human reality must this graph represent?
Organizational Intent	What outcome should the assistant support, and what should it avoid?
Knowledge & Data	What must the assistant know, connect, retrieve, or distinguish from inference?
Governance & Trust	What must this graph prevent, caveat, refuse, or escalate?
Employee Enablement	Who would maintain, correct, review, or recover this part of the experience?
Interaction & Service Delivery	How does the guest feel this graph in the moment?
Human-AI Learning & Adaptation	What should improve when users correct, challenge, or abandon guidance?

Core entities

- Option
- Recommendation
- Decision State
- Decision Hierarchy
- Decision Criterion
- Tradeoff
- Constraint
- Priority
- Comparison
- Explanation
- Next Best Step
- Navigation Cue
- Alternative Path
- Recovery Option
- Signal
- Noise
- Friction Point

- Confidence Prompt
- Confidence Gap
- User Correction

Core relationships

- Recommendation includes Option
- Option has Tradeoff
- Tradeoff relates to Decision Criterion
- Decision Criterion belongs to Decision Hierarchy
- Constraint changes Next Best Step
- Friction Point creates Confidence Gap
- Confidence Gap triggers Confidence Prompt
- User Correction updates Priority
- Signal supports Recommendation
- Noise should be suppressed
- Alternative Path resolves Conflict
- Recovery Option resolves Disruption
- Explanation clarifies Tradeoff
- Navigation Cue supports Wayfinding

Guardrails

- Do not provide a flat list when the user needs decision structure.
- Do not hide the reason for a recommendation.
- Do not overwhelm the user with true but irrelevant facts.
- Do not treat user correction as failure; use it as input.
- Do not assume one objectively best answer when priorities are in tension.

Responsible behaviors this graph should enable

- Explain why a recommendation fits.
- Name tradeoffs clearly.
- Separate signal from noise.
- Simplify options when the user is overwhelmed.
- Use corrections to update priorities.
- Offer confidence prompts such as easiest, fastest, most fun, lower walking, or included-only.

Sample assistant behavior

Because lunch is in 35 minutes, I would make reservation protection the top priority. Your best options are: stay nearby for a short attraction if the wait is low, take a shaded break, or check in early. I would not cross the park right now unless you are comfortable with a tight buffer.

QA prompts

1. Space Mountain is closed. What should we do now?
2. We have lunch in 35 minutes. What is the best next step?
3. That recommendation is too far.
4. I am overwhelmed. Just tell me what to do.
5. Which is better: Pirates now, Haunted Mansion now, or taking a break?

Pass criteria

- The answer follows the graph purpose, not generic Disneyland advice.
- The answer asks or caveats when the system lacks enough context.
- The answer preserves human agency and creates a human-control moment when needed.
- The answer avoids unsupported certainty, hidden steering, and sensitive inference.
- The answer uses language that fits the risk of the moment.