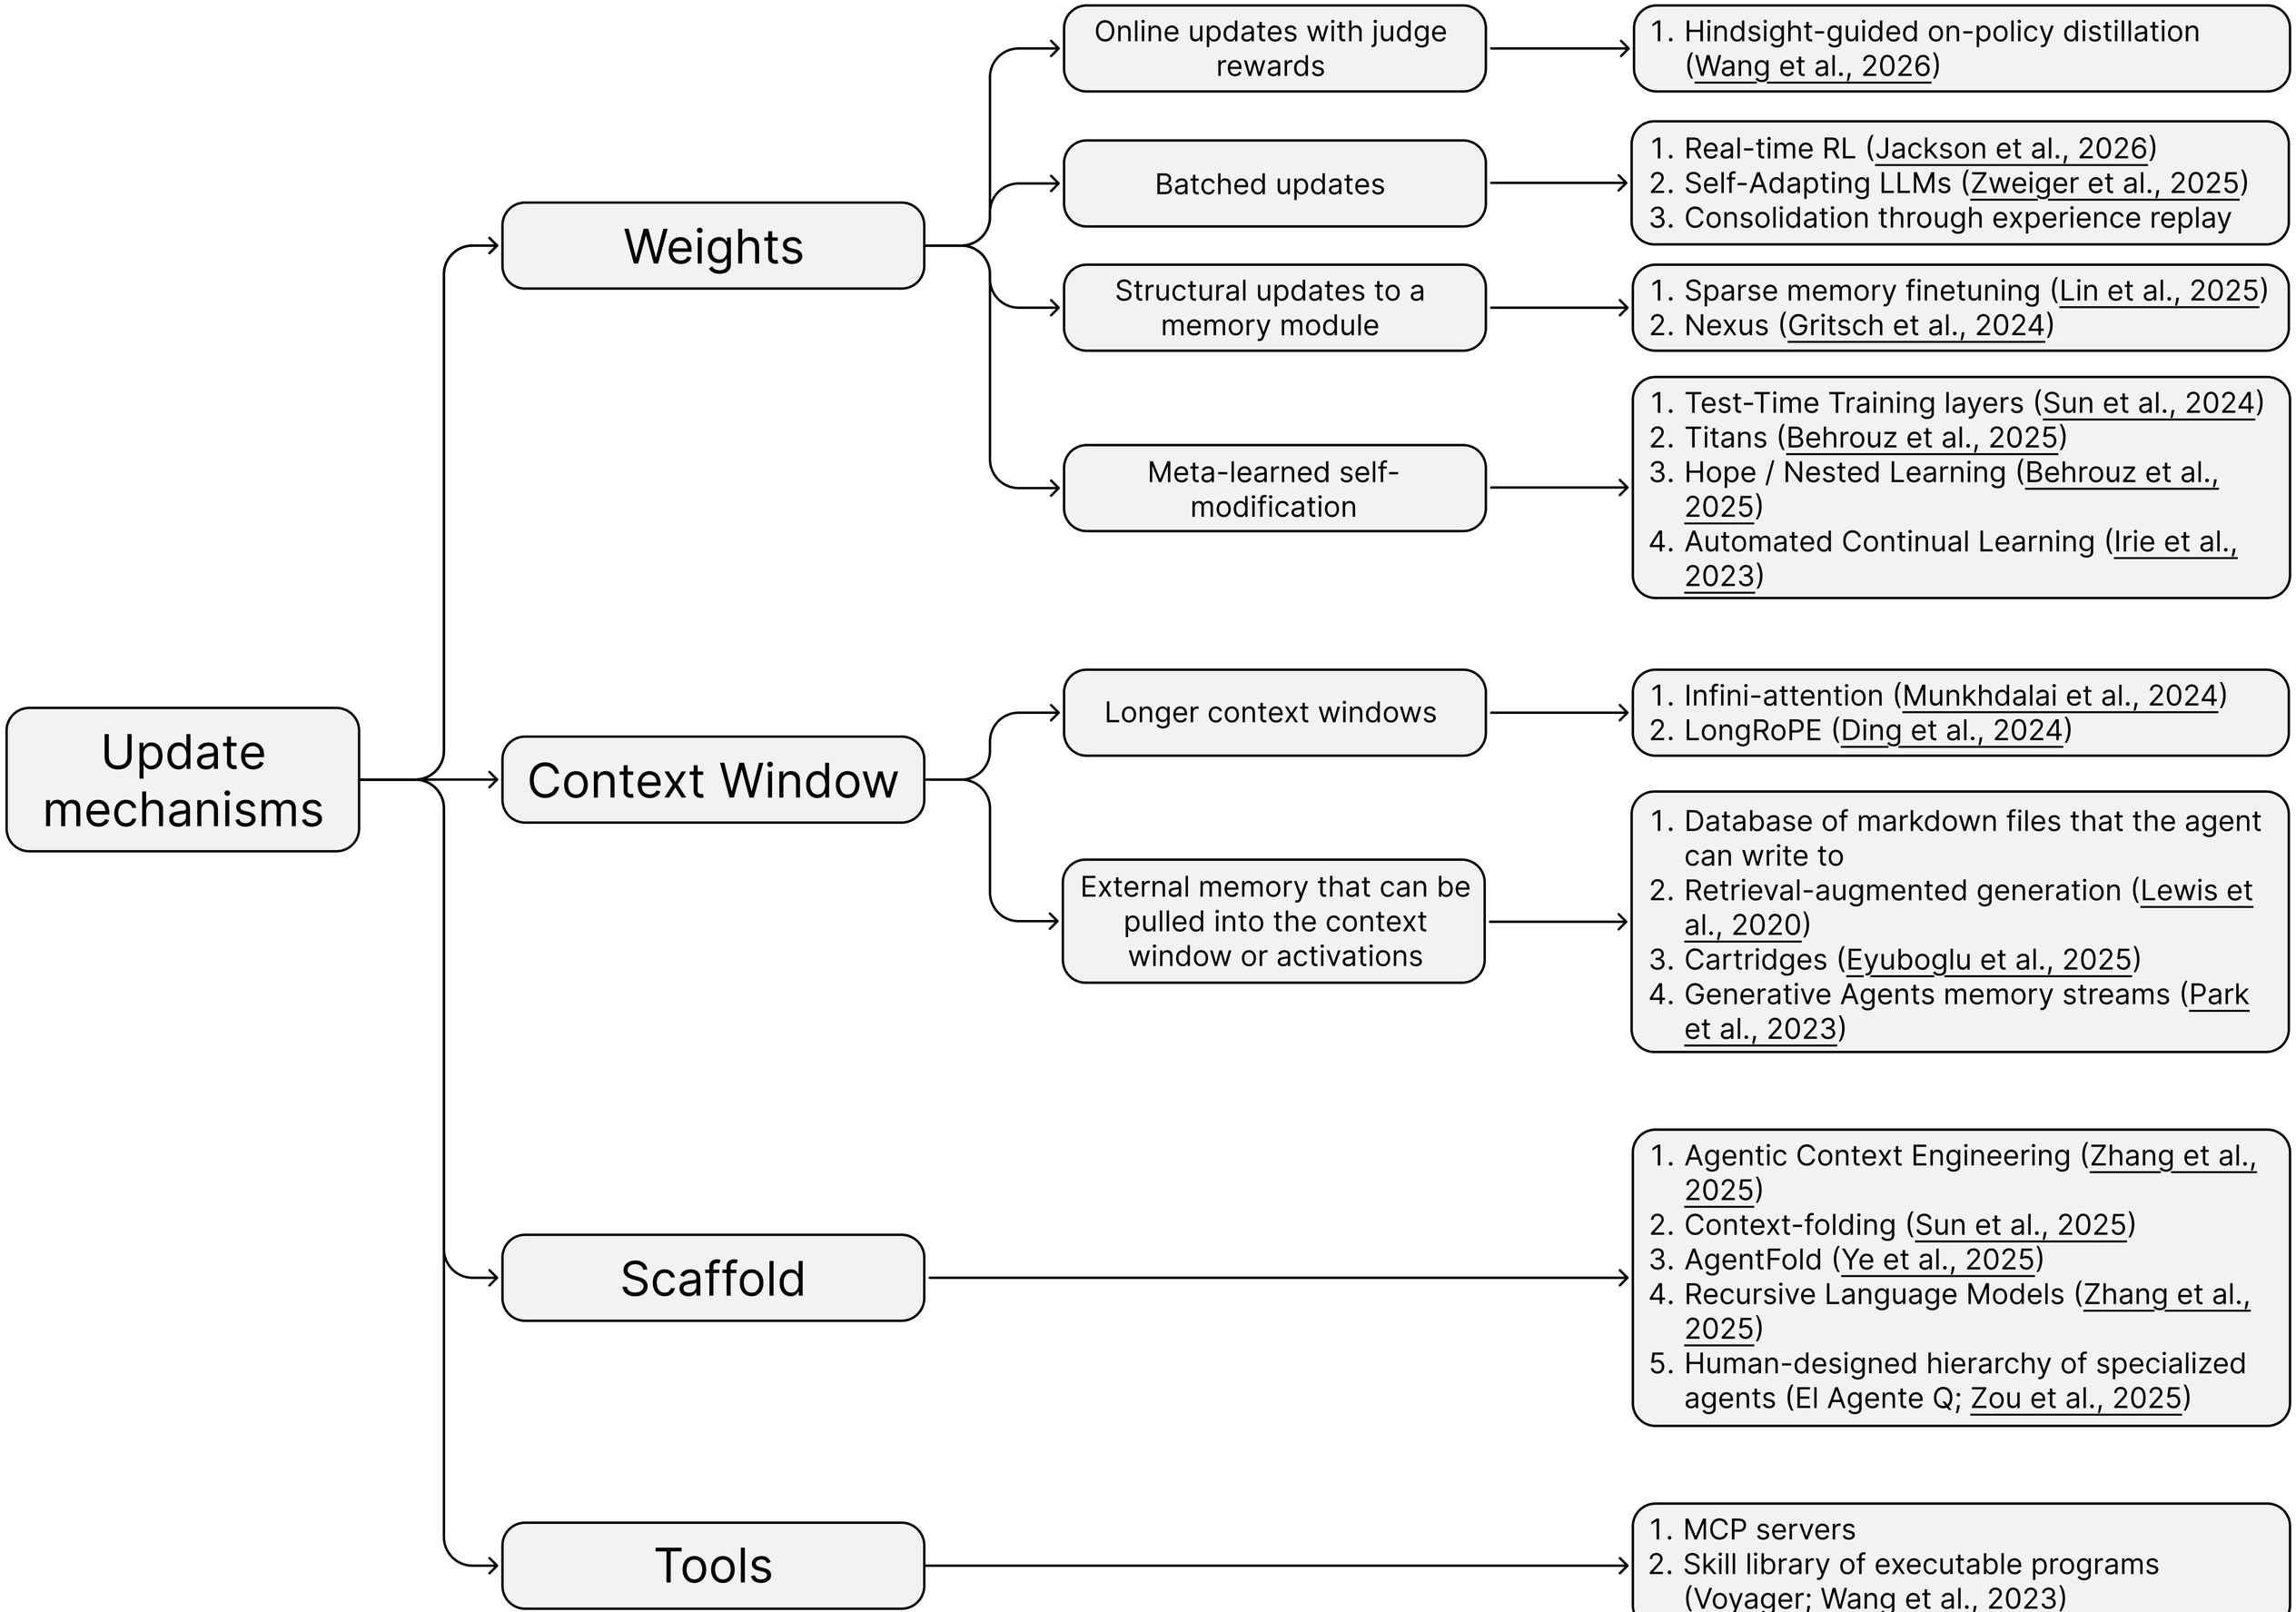
Update mechanisms
Weights
Online updates with judge rewards
1. Hindsight-guided on-policy distillation (Wang et al., 2026)
Batched updates
1. Real-time RL (Jackson et al., 2026)
2. Self-Adapting LLMs (Zweiger et al., 2025)
3. Consolidation through experience replay
Structural updates to a memory module
1. Sparse memory finetuning (Lin et al., 2025)
2. Nexus (Gritsch et al., 2024)
Meta-learned self-modification
1. Test-Time Training layers (Sun et al., 2024)
2. Titans (Behrouz et al., 2025)
3. Hope / Nested Learning (Behrouz et al., 2025)
4. Automated Continual Learning (Irie et al., 2023)
Context Window
Longer context windows
1. Infini-attention (Munkhdalai et al., 2024)
2. LongRoPE (Ding et al., 2024)
External memory that can be pulled into the context window or activations
1. Database of markdown files that the agent can write to
2. Retrieval-augmented generation (Lewis et al., 2020)
3. Cartridges (Eyuboglu et al., 2025)
4. Generative Agents memory streams (Park et al., 2023)
Scaffold
1. Agentic Context Engineering (Zhang et al., 2025)
2. Context-folding (Sun et al., 2025)
3. AgentFold (Ye et al., 2025)
4. Recursive Language Models (Zhang et al., 2025)
5. Human-designed hierarchy of specialized agents (El Agente Q; Zou et al., 2025)
Tools
1. MCP servers
2. Skill library of executable programs (Voyager; Wang et al., 2023)