

Supplementary Information for

Evidence-grounded AI for longitudinal clinical management in musculoskeletal care

Wenjie Li^{1,2,3,†}, Yujie Zhang^{2,3,4,†}, Fanrui Zhang^{2,5,†}, Haoran Sun^{3,6,†}, Renhao Yang^{1,†}, Junjun He^{2,3}, Weiran Huang², Yuanfeng Ji⁷, Chenrun Wang², Kailing Wang⁸, Hongcheng Gao⁹, Kaipeng Zhang², Hanyu Wang¹, Angela Lin Wang¹, Xingqi He⁴, Yilin Huang¹, Shiyi Yao¹, Lilong Wang³, Yankai Jiang³, Yirong Chen³, Chenglong Ma^{2,3}, Jiyao Liu³, Ming Hu³, Gen Li¹, Yidong Xu¹, Chengyu Zhuang¹, Jiawei Liu⁵, Yin Zhang¹, Lequan Yu⁸, Lu Chen², Yinpeng Dong⁹, Lei Liu⁶, Carlos Gutiérrez Sanromán¹⁰, Yu Qiao^{2,3,*}, Weijie Ma^{2,*}, Xiaosong Wang^{2,3,*}, Lei Wang^{1,*}

¹Department of Orthopaedics, Ruijin Hospital, College of Health Science and Technology, Shanghai Jiao Tong University School of Medicine, Shanghai, China.

²Shanghai Innovation Institute, Shanghai, China.

³Shanghai Artificial Intelligence Laboratory, Shanghai, China.

⁴College of Computer Science and Artificial Intelligence, Fudan University, Shanghai, China.

⁵MoE Key Laboratory of Brain-inspired Intelligent Perception and Cognition, University of Science and Technology of China, Hefei, China.

⁶School of Basic Medical Sciences, Intelligent Medicine Institute, Fudan University, Shanghai, China.

⁷Department of Radiation Oncology, Stanford University School of Medicine, Stanford, USA.

⁸School of Computing and Data Science, The University of Hong Kong, Hong Kong SAR, China.

⁹College of AI, Tsinghua University, Beijing, China.

¹⁰Servicio Cirugia Pediátrica, Department of Pediatric Surgery, Hospital Universitario y Politécnico La Fe, Valencia, Spain.

*Corresponding author(s). E-mail(s): qiaoyu@pjlab.org.cn; weijiema@sii.edu.cn; wangxiaosong@pjlab.org.cn; wl10448@rjh.com.cn;

[†]These authors contributed equally to this work.

This PDF file includes

Supplementary Methods

Supplementary Results

Supplementary References

1324

Contents

1325

Evidence-grounded AI for longitudinal clinical management in musculoskeletal care

1326

1327

A Main Workflow of OrthoPilot	3
--------------------------------------	----------

1328

B Complete Task-Level Results on OrthoBench	3
--	----------

1329

C Centre-Level Results of External Multicentre Validation	4
--	----------

1330

D Ablation Study	6
-------------------------	----------

1331

D.1 Analysis on Models for Agentic System	6
---	---

1332

D.2 Analysis on Agentic Components	6
--	---

1333

D.3 Analysis on Progressive Training Strategy	7
---	---

1334

D.4 Analysis on LLM Backbones	7
---	---

1335

D.5 Analysis on Model Parameters	8
--	---

1336

E Details of Datasets	8
------------------------------	----------

1337

E.1 Retrospective Dataset Curation	8
--	---

1338

E.2 Prospective Dataset Curation	8
--	---

1339

E.3 Synthetic Dataset Curation	9
--	---

1340

F Implementation of CHEESE	12
-----------------------------------	-----------

1341

F.1 Details of Backbone	12
-----------------------------------	----

1342

F.2 Implementation of Training Details	12
--	----

1343

G Baselines	13
--------------------	-----------

1344

G.1 Proprietary Models	13
----------------------------------	----

1345

G.2 Open-Source General-Purpose Models	13
--	----

1346

G.3 Medical-Specific Models	14
---------------------------------------	----

1347

G.4 Agentic Framework	14
---------------------------------	----

1348

H OrthoBench Design and Evaluation Metrics	14
---	-----------

1349

H.1 Task Design	17
---------------------------	----

1350

H.2 Data Composition	18
--------------------------------	----

1351

H.3 Physician Evaluation of ORACLE	19
--	----

1352

I Details of Reader Study	22
----------------------------------	-----------

1353

J Rating Scales and Questionnaires	26
---	-----------

1354

K Case Study	35
---------------------	-----------

1355

K.1 Full-Pathway Cases	35
----------------------------------	----

1356

K.2 OrthoBench Cases	42
--------------------------------	----

1357

K.3 Application Cases	42
---------------------------------	----

1358

L Prompts	60
------------------	-----------

1359

L.1 Agent Prompts	60
-----------------------------	----

1360

L.2 QA generation prompt	60
------------------------------------	----

1361

L.3 Training prompt	61
-------------------------------	----

1362

L.4 Evaluation prompts	61
----------------------------------	----

1363

Supplementary References	97
---------------------------------	-----------

1364 Appendix A Main Workflow of OrthoPilot

1365 The single-query workflow of OrthoPilot is summarized in Algorithm 1. The main agent first analyzes the clinical query
 1366 and the current patient state. It then selects actions from Tool Plaza, dispatches specialized sub-agents, and collects
 1367 returned reports from the in-hospital information world, the external knowledge world, or the multimodal interface.
 1368 The final output is produced after evidence synthesis.

Algorithm 1 Main Workflow of OrthoPilot for a Single Clinical Query

```

1: Input: clinical query  $q$ ; patient state  $s$ ; Tool Plaza  $\mathcal{T} = \mathcal{T}_{\text{in}} \cup \mathcal{T}_{\text{ext}} \cup \mathcal{T}_{\text{mm}}$ 
2: Symbols:  $\tau_t$  reasoning trace;  $a_t$  selected action;  $o_t$  returned report;  $H_t$  trajectory history
3: Output: structured clinical recommendation  $y$ 
4: Initialize  $\mathcal{X}_{\text{in}} \leftarrow \emptyset$ ,  $\mathcal{K}_{\text{ext}} \leftarrow \emptyset$ ,  $H_0 \leftarrow \emptyset$ 

5: for  $t = 0, 1, \dots, T$  do
6:    $\tau_t \leftarrow \text{Reason}_\theta(q, s, \mathcal{X}_{\text{in}}, \mathcal{K}_{\text{ext}}, H_t)$  ▷ main-agent reasoning
7:   if  $\text{Ready}(\tau_t)$  then
8:     break ▷ evidence is sufficient
9:   end if
10:   $a_t \leftarrow \text{Act}(\tau_t, \mathcal{T})$  ▷ choose next action
11:  if  $a_t \in \mathcal{T}_{\text{in}}$  then
12:    dispatch in-hospital sub-agent with query  $q_t^{\text{in}}$ 
13:     $o_t \subseteq \bigcup_{n=1}^N \mathcal{X}_n(q_t^{\text{in}})$  ▷ return patient evidence
14:     $\mathcal{X}_{\text{in}} \leftarrow \mathcal{X}_{\text{in}} \cup o_t$ 
15:  else if  $a_t \in \mathcal{T}_{\text{ext}}$  then
16:    dispatch external-knowledge sub-agent with query  $q_t^{\text{ext}}$ 
17:     $o_t \subseteq \bigcup_{m=1}^M \mathcal{K}_m(q_t^{\text{ext}})$  ▷ return external evidence
18:     $\mathcal{K}_{\text{ext}} \leftarrow \mathcal{K}_{\text{ext}} \cup o_t$ 
19:  else
20:    dispatch multimodal sub-agent with query  $q_t^{\text{mm}}$ 
21:     $o_t \leftarrow \text{Parse}_{\text{mm}}(q_t^{\text{mm}}, s)$  ▷ return multimodal report
22:     $\mathcal{X}_{\text{in}} \leftarrow \mathcal{X}_{\text{in}} \cup o_t$ 
23:  end if
24:   $H_{t+1} \leftarrow (\tau_0, a_0, o_0, \dots, \tau_t, a_t, o_t)$  ▷ update ReAct trajectory
25: end for

26:  $y \leftarrow \text{Synthesize}_\theta(q, s, \mathcal{X}_{\text{in}}, \mathcal{K}_{\text{ext}}, H_T)$  ▷ final synthesis
27: return  $y$ 

```

1369 Appendix B Complete Task-Level Results on OrthoBench

Table B1: Complete task-level performance across the 11 OrthoBench tasks. Values are percentages. Standalone baselines and OrthoPilot variants are shown together for direct comparison. Best and second-best values in each column are shown in bold and underlined, respectively.

Model	Closed-ended diagnostic tasks				Open-ended management tasks							Averages	
	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	T11	Diag. avg	Mgmt. avg
Reasoning LLMs													
Grok-4	53.67	54.10	66.91	61.06	43.65	62.56	72.90	66.21	82.88	86.80	70.10	58.94	69.30
GPT-5.1	56.60	52.67	66.36	60.56	50.21	62.71	75.80	57.13	78.26	84.77	70.40	59.05	68.47
Kimi-K2	55.53	49.90	63.77	59.29	41.87	56.22	72.35	60.26	82.43	77.08	69.90	57.12	65.73
DeepSeek-R1	52.37	48.73	65.36	60.53	36.40	47.16	69.68	61.41	82.60	83.80	69.90	56.75	64.42
Qwen3-32B	52.60	45.83	59.93	55.79	35.31	44.28	57.80	56.18	74.09	72.73	65.61	53.54	58.00
GPT-oss-120B	51.80	51.07	63.98	60.59	43.46	27.49	53.83	60.04	83.80	83.59	69.41	56.86	60.23
Claude-Sonnet-4.5	57.50	53.50	68.00	61.00	43.73	50.63	72.80	54.13	76.28	77.00	70.00	60.00	63.51

Continued on next page

Table B1 (continued)

Model	Closed-ended diagnostic tasks				Open-ended management tasks							Averages	
	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	T11	Diag. avg	Mgmt. avg
General LLMs													
Gemini-2.5-flash	52.33	49.63	65.13	60.39	29.60	62.74	57.16	24.26	46.60	70.83	66.09	56.87	51.04
DeepSeek-V3	52.83	46.83	64.76	59.39	36.68	56.27	72.15	67.44	68.32	69.36	58.18	55.95	61.20
Mistral-Large	47.96	45.55	57.75	52.74	36.89	50.52	64.82	60.58	74.66	71.55	65.17	51.00	60.60
Qwen3-235B-A22B	52.90	50.67	61.46	56.59	34.28	41.25	62.78	60.09	81.77	86.58	69.90	55.40	62.38
GPT-4o	51.03	50.33	61.18	55.99	20.38	39.04	51.40	47.35	74.45	56.28	67.47	54.63	50.91
Llama-4-Maverick	49.93	46.23	65.44	58.39	26.58	27.69	43.08	46.75	68.24	56.98	49.61	55.00	45.56
Phi-4	34.37	35.65	37.69	35.67	23.02	26.19	23.94	43.40	69.45	58.49	53.50	35.84	42.57
Llama-3.3-70B	46.90	47.57	58.77	53.62	18.84	26.00	29.52	42.76	64.11	54.83	56.95	51.72	41.86
Medical LLMs													
MedGemma-27B	47.88	46.35	57.25	51.82	38.84	56.61	59.60	65.22	75.12	84.04	63.18	50.82	63.23
HuatuoGPT-o1-72B	48.30	48.23	56.11	52.52	32.36	23.13	37.62	37.27	70.17	49.86	55.36	51.29	43.68
OpenBioLLM-70B	42.57	44.40	49.76	45.49	17.85	22.47	31.89	29.16	61.58	42.69	49.71	45.56	36.48
Agent methods													
MDAgents (DeepSeek-V3)	35.04	47.32	56.88	38.76	41.55	37.79	51.85	45.25	36.46	39.22	41.87	44.50	42.00
MDAgents (GPT-4o)	47.53	35.48	46.25	38.74	43.58	43.10	42.10	42.67	45.53	38.82	27.70	42.00	40.50
Ours													
CHEESE OrthoPilot (CHEESE)	63.82	54.00	68.31	61.87	48.40	62.80	73.00	70.00	84.00	87.00	71.80	62.00	71.00
	84.20	78.20	83.60	81.80	83.10	78.00	85.50	86.00	88.50	89.40	85.40	81.95	85.13
OrthoPilot (GPT-5.1)	<u>79.95</u>	<u>75.01</u>	<u>81.61</u>	<u>79.99</u>	<u>75.78</u>	<u>76.63</u>	<u>84.55</u>	80.49	87.72	<u>90.87</u>	<u>83.93</u>	<u>79.14</u>	<u>82.85</u>
OrthoPilot (DeepSeek-R1)	76.09	71.62	79.58	78.88	67.41	69.31	80.91	79.28	<u>88.46</u>	90.93	82.08	76.54	79.77
OrthoPilot (GPT-4o)	76.39	72.59	78.93	77.52	71.13	71.54	79.27	79.54	87.67	84.66	83.83	76.36	79.66
OrthoPilot (MedGemma)	73.03	69.18	76.00	73.97	66.85	72.21	77.32	<u>80.58</u>	85.00	89.97	79.19	73.05	78.73

T1, admission diagnosis; T2, preoperative diagnosis; T3, intraoperative diagnosis; T4, discharge diagnosis; T5, perioperative assessment; T6, surgical planning; T7, preoperative orders; T8, postoperative orders; T9, discharge summary; T10, rehabilitation planning; T11, multidisciplinary consultation. Diag. avg, mean across T1–T4. Mgmt. avg, mean across T5–T11. OrthoPilot (X) denotes the OrthoPilot framework instantiated with backbone X. OrthoPilot (CHEESE) denotes the final full system.

1370

Appendix C Centre-Level Results of External Multicentre Validation

Table C2: Centre-level external multicentre validation results. Values are percentages. Results are shown for closed-ended diagnostic tasks and open-ended management tasks across 60 external hospital centres. Within each centre, the highest closed-ended score is shown in bold with a blue shaded cell, and the highest open-ended score is shown in bold with a orange shaded cell.

Centre	n	Closed-ended diagnostic tasks				Open-ended management tasks			
		Ortho Pilot	DeepSeek -R1	Qwen3 -235B	MedGemma -27B	Ortho Pilot	DeepSeek -R1	Qwen3 -235B	MedGemma -27B
Cohort 01	112	70.85	66.58	51.09	59.25	79.32	68.88	56.88	65.65
Cohort 02	125	72.48	52.90	55.68	48.07	70.76	74.00	63.26	68.42
Cohort 03	117	70.11	70.84	59.61	51.92	70.28	65.77	63.11	53.46
Cohort 04	150	76.69	53.78	71.89	61.50	91.12	73.80	76.34	80.40
Cohort 05	144	76.01	57.31	43.92	49.03	72.33	61.46	60.35	81.10
Cohort 06	144	79.32	68.58	54.34	56.82	81.25	77.81	66.31	67.95

Continued on next page

Table C2 (continued)

Centre	n	Closed-ended diagnostic tasks				Open-ended management tasks			
		Ortho Pilot	DeepSeek -R1	Qwen3 -235B	MedGemma -27B	Ortho Pilot	DeepSeek -R1	Qwen3 -235B	MedGemma -27B
Cohort 07	111	79.98	57.15	64.48	60.05	64.34	63.00	69.87	69.15
Cohort 08	109	85.75	57.31	55.22	53.71	96.04	86.56	71.93	73.81
Cohort 09	150	82.17	53.51	62.31	44.89	76.26	64.14	62.84	64.49
Cohort 10	147	67.48	53.79	63.33	45.51	91.36	51.09	63.07	53.43
Cohort 11	125	69.60	50.13	48.23	42.45	68.50	69.22	57.61	64.73
Cohort 12	38	82.29	54.99	56.07	54.16	79.52	53.48	58.20	44.84
Cohort 13	136	76.45	45.83	70.56	57.66	75.22	67.91	62.09	65.01
Cohort 14	109	76.58	58.36	56.32	32.73	94.29	67.54	67.84	74.44
Cohort 15	123	45.97	35.35	39.66	37.06	80.98	73.09	74.76	67.79
Cohort 16	28	77.79	59.57	42.84	51.81	81.30	48.33	62.89	71.51
Cohort 17	107	66.45	42.89	61.13	36.34	70.49	57.55	66.24	44.40
Cohort 18	145	72.98	54.89	57.95	50.82	82.90	52.33	49.57	56.41
Cohort 19	94	67.68	57.88	51.97	48.42	70.41	47.91	67.43	63.13
Cohort 20	144	46.53	45.80	63.94	37.76	85.77	71.85	65.99	70.27
Cohort 21	76	75.52	65.41	54.94	55.65	86.00	72.47	68.56	69.03
Cohort 22	135	79.85	64.26	50.59	58.66	66.98	60.46	32.43	52.59
Cohort 23	118	81.79	50.28	31.05	59.84	84.70	65.80	77.93	82.96
Cohort 24	86	59.08	51.68	40.20	49.39	81.87	67.61	65.86	74.38
Cohort 25	24	64.89	54.69	51.07	38.74	79.04	61.85	66.60	60.94
Cohort 26	135	63.92	62.78	45.99	45.31	62.30	52.28	47.66	68.43
Cohort 27	131	71.07	42.06	58.16	47.68	77.54	78.96	68.52	75.30
Cohort 28	68	77.93	50.87	47.35	56.69	80.64	78.03	57.24	53.05
Cohort 29	79	71.20	55.62	48.39	50.58	96.70	65.20	58.16	75.29
Cohort 30	128	82.26	60.61	48.94	69.34	90.13	76.98	84.64	80.00
Cohort 31	149	73.21	55.52	63.59	45.38	77.15	70.96	62.22	60.80
Cohort 32	129	92.77	72.33	52.49	68.80	89.58	44.30	65.11	57.14
Cohort 33	141	70.52	50.89	61.73	59.32	59.33	58.97	62.21	57.90
Cohort 34	122	63.27	54.52	53.13	53.94	64.12	59.04	62.80	60.77
Cohort 35	150	68.98	67.94	51.27	52.80	80.46	57.56	71.53	64.52
Cohort 36	122	58.85	47.29	59.40	48.09	78.23	75.90	72.51	65.98
Cohort 37	131	79.59	70.53	62.30	48.27	68.34	74.55	74.67	62.84
Cohort 38	132	61.01	52.80	41.86	34.38	83.21	70.03	64.94	69.51
Cohort 39	133	64.56	38.61	52.95	40.43	83.60	61.91	67.01	60.35
Cohort 40	136	74.45	70.10	54.17	52.41	86.26	80.72	79.26	73.80
Cohort 41	72	64.99	51.15	61.28	70.74	80.93	86.64	85.72	80.73
Cohort 42	67	73.22	61.63	59.71	33.44	83.00	63.81	73.35	40.24
Cohort 43	132	75.64	60.44	67.86	55.46	79.28	63.94	69.31	65.62
Cohort 44	70	66.89	49.80	46.65	35.53	83.69	63.42	56.30	57.07
Cohort 45	124	66.34	55.08	57.33	42.65	92.44	64.08	75.31	70.91
Cohort 46	79	75.86	51.70	36.96	55.38	79.28	80.09	60.02	64.10
Cohort 47	131	85.54	60.33	45.74	49.07	70.05	72.50	63.06	56.25
Cohort 48	114	82.24	75.69	64.49	46.22	73.08	67.75	56.51	64.38
Cohort 49	142	78.48	58.38	60.59	51.89	79.88	73.26	62.62	65.79
Cohort 50	145	59.97	43.61	39.23	46.70	86.74	68.08	67.81	71.42
Cohort 51	147	56.46	54.52	55.42	55.64	74.24	78.51	74.92	64.49
Cohort 52	116	73.26	56.70	51.57	53.21	81.51	65.66	69.75	66.77
Cohort 53	115	71.53	51.73	42.51	51.26	73.67	80.41	88.87	68.70
Cohort 54	135	94.05	61.58	67.03	50.99	91.85	64.54	77.95	62.28
Cohort 55	99	80.29	59.29	58.15	71.89	83.18	86.53	70.34	73.93
Cohort 56	27	84.49	67.34	54.61	61.75	85.45	73.18	73.00	65.81
Cohort 57	33	67.67	53.39	58.71	29.45	75.64	64.05	71.19	60.28
Cohort 58	113	63.23	67.02	55.79	51.84	79.55	63.51	58.36	59.02
Cohort 59	116	89.17	62.38	60.66	45.90	81.84	67.21	74.01	79.03
Cohort 60	136	78.99	80.77	67.46	65.47	73.03	67.87	65.83	70.05

Closed-ended diagnostic tasks are evaluated by accuracy. Open-ended management tasks are evaluated by the Open Response Assessment for Clinical Language Evaluation framework (ORACLE) score. Blue shading marks the highest closed-ended score within each centre. Orange shading marks the highest open-ended score within each centre. All values are centre-level percentages.

Appendix D Ablation Study

D.1 Analysis on Models for Agentic System

We next asked whether OrthoPilot depended on CHEESE, or whether the same workflow could strengthen other frontier Large Language Models (LLMs). We therefore kept the agent architecture fixed and replaced only the central reasoning model with DeepSeek-R1, GPT-4o, MedGemma-27B and GPT-5.1 on OrthoBench.

The same pattern was preserved across backbones (Fig. D1a). CHEESE gave the strongest instantiation of the framework, with an overall score of 83.97%. GPT-5.1 was the next strongest backbone, reaching 81.50% overall. The other models were lower, but they followed the same performance structure across diagnostic and management tasks. Thus, the benefit of OrthoPilot was not tied to a single host model.

The paired comparison in Fig. D1b reached the same conclusion. OrthoPilot improved every tested backbone over its standalone counterpart, with overall gains ranging from 16.25% to 26.20%. This supports the transferability of the shared agent workflow. It also points to the core mechanism of the system. Full-pathway musculoskeletal care requires repeated evidence gathering, state tracking and decision revision. The agent layer supplies this structure across LLM families, whereas CHEESE defines the highest ceiling within that workflow.




Fig. D1: Ablation analysis of the OrthoPilot agent framework and tool system. **a**, OrthoBench performance after replacing the backbone LLM in the fixed OrthoPilot workflow. **b**, Paired comparison between each standalone backbone and its OrthoPilot agent counterpart, showing that the workflow transfers across LLM families. **c**, Tool-system ablation of full OrthoPilot, variants without in-hospital or external tools, and CHEESE standalone. Together, these analyses show that continuous clinical decision-making depends on both a transferable agent workflow and access to patient-grounded and external evidence.

D.2 Analysis on Agentic Components

We then tested which evidence tools were most important for continuous clinical decision-making. We disabled the in-hospital tool system and the external tool system separately while preserving the remaining OrthoPilot architecture, and compared both variants with full OrthoPilot and CHEESE standalone.

Removing in-hospital tools produced the largest loss (Fig. D1c). Overall performance fell from 83.97% to 73.37%, a drop of 10.60%. This identifies patient-specific evidence as the main substrate for full-pathway reasoning. OrthoPilot must track records, imaging, laboratory data and consultations as the patient state changes. Without this stream, the system loses much of its ability to revise decisions over time.

Removing external tools caused a smaller but still clear decline, with overall performance falling to 76.98%. CHEESE standalone was lower still, at 67.73% overall. These results show that external knowledge is not redundant. It calibrates patient-level reasoning against guidelines, literature and biomedical knowledge. The two evidence




Fig. D2: Comprehensive ablation study of CHEESE. **a**, Stage-wise ablation of the progressive training curriculum, showing CHEESE performance on diagnostic tasks (accuracy, blue) and management tasks (ORACLE score, orange) when each training stage is individually removed. **b**, Impact of LLM backbone on CHEESE clinical reasoning at identical scale (8B parameters), comparing InternLM3, LLaMA-3, and Qwen3 under the same training curriculum. **c**, Model-size scaling behaviour of CHEESE across Qwen3 backbone variants from 4B to 32B parameters.

worlds therefore play complementary roles. Patient data anchors the evolving clinical state. External knowledge keeps decisions aligned with current practice.

D.3 Analysis on Progressive Training Strategy

To quantify the contribution of each curriculum stage to CHEESE, we conducted a stage-wise ablation in which a single stage was removed while the remaining four were retained. Performance was evaluated on OrthoBench using mean accuracy for closed-ended diagnostic tasks (Tasks 1–4) and mean ORACLE score for open-ended management tasks (Tasks 5–11). As shown in Fig. D2a, the five stages contributed unevenly, with a clear hierarchy of importance among the tool-free stages. Removing Stage II (orthopaedic reasoning cold-start alignment) caused the largest decline, lowering diagnostic accuracy from 62.00% to 55.18% and management score from 71.00% to 61.07%. This identifies task-specific supervised fine-tuning on the 21-year longitudinal cohort as the decisive step for internalising the domain-specific decision structures spanning all 11 pathway tasks. Stage III (reason-driven reinforcement learning with Group Relative Policy Optimization (GRPO)) ranked second; its removal reduced diagnostic accuracy to 57.36% and management score to 64.52%, reflecting the role of reinforcement learning in sharpening reasoning quality, structural compliance and self-calibration on ambiguous cases. Stage I (clinical domain adaptation) produced the smallest but still material decline of the three, with diagnostic accuracy of 59.18% and management score of 67.45%, indicating that broad medical knowledge from large-scale domain pre-training establishes the ceiling of subsequent specialisation even though part of it can be recovered downstream. Importantly, all three ablated variants remained above the untrained Qwen3-32B baseline of 53.54% / 58.00%, confirming that any single stage already delivers measurable gains over the raw foundation model, while their cumulative composition is what unlocks full CHEESE performance.

By contrast, Stages IV and V, which operate in the multi-turn and tool-augmented regime, had only marginal influence on standalone CHEESE performance. Removing Stage IV (multi-horizon clinical trajectory distillation) yielded 60.47% / 69.82% and removing Stage V (full-pathway agentic reinforcement learning) yielded 61.03% / 70.18%. This pattern is consistent with the architectural expectation that trajectory distillation and agentic reinforcement learning primarily benefit the complete OrthoPilot system, where CHEESE interacts with external tools, rather than the standalone reasoning engine evaluated without tool access. Together, these results validate the progressive curriculum: Stages I–III establish a strong tool-free clinical reasoning core through domain adaptation, task alignment and reinforcement learning, while Stages IV–V equip the model with multi-horizon trajectory reasoning and agentic tool-use capabilities that are fully realised only within the complete OrthoPilot orchestration framework.

D.4 Analysis on LLM Backbones

To isolate the effect of backbone architecture and pre-training corpus on clinical performance, we trained three CHEESE variants with backbones of identical scale (8B parameters) drawn from different model families—InternLM3-8B, LLaMA-3-8B and Qwen3-8B—under the same five-stage curriculum, hyperparameter configuration and data mixture. As shown in Fig. D2b, backbone selection substantially affected both task families. Qwen3-8B achieved the best results, reaching 54.13% diagnostic accuracy and 62.07% management score, followed by InternLM3-8B at 49.72% / 57.86% and LLaMA-3-8B at 47.58% / 55.94%. These gaps likely reflect both architectural design and pre-training composition: Qwen3 incorporates a broader multilingual and scientific corpus that provides a stronger initialisation for medical domain adaptation, whereas the predominantly English-centric corpus of LLaMA-3 limits its capacity to encode Chinese-language clinical narratives, which constitute the majority of the Ruijin Hospital longitudinal cohort. Based on these findings, Qwen3 was selected as the backbone for all subsequent CHEESE training and deployment.

1435 D.5 Analysis on Model Parameters

1436 To characterise the relationship between model capacity and clinical reasoning performance, we trained the full
1437 CHEESE curriculum on four Qwen3 backbone variants spanning an eight-fold parameter range—Qwen3-4B, Qwen3-
1438 8B, Qwen3-14B and Qwen3-32B—under identical data mixtures and training stages, with batch size and learning rate
1439 adjusted to accommodate differing capacities. As shown in Fig. D2c, both diagnostic accuracy and management per-
1440 formance improved monotonically with model scale. Qwen3-32B led both task families with 62.00% / 71.00%, followed
1441 by Qwen3-14B at 57.91% / 66.24%, Qwen3-8B at 54.13% / 62.07% and Qwen3-4B at 44.86% / 52.33%. Scaling from
1442 4B to 32B thus raised diagnostic accuracy from 44.86% to 62.00% and management score from 52.33% to 71.00%,
1443 establishing model capacity as a primary determinant of clinical reasoning quality.

1444 Appendix E Details of Datasets

1445 E.1 Retrospective Dataset Curation

1446 In this section, we describe how the raw clinical cohorts were curated and allocated for benchmark construction and
1447 retrospective validation. The full workflow is summarized in Fig. E3. Within the Ruijin cohort, all benchmark partitions
1448 were defined at the patient level, so that each patient was assigned once and all downstream task records from that
1449 patient remained in the same partition.

1450 We first curated the internal Ruijin cohort to construct OrthoBench. Starting from 183,472 patients in the longitu-
1451 dinal orthopaedic cohort spanning 2004 to 2024, we reserved earlier in-house cases from 2004 to 2011 for the historical
1452 similar-case retrieval resource ($n = 48,464$). This resource was defined by calendar time and used only for retrieval.
1453 It was kept separate from benchmark curation. We then restricted the remaining cohort to patients with extractable
1454 discharge diagnosis coding and sufficient disease-level information for stratified sampling, yielding 89,746 patients
1455 for dataset construction. Patient-level allocation yielded OrthoBench, with a fixed test cohort of 5,905 patients. The
1456 remaining 83,841 patients were used for construction of the training cohort. We preserved broad disease diversity and
1457 retained 1,000 disease categories in OrthoBench.

1458 We next curated an independent multicentre cohort for retrospective validation. From 60 participating hospitals,
1459 we assembled 8,018 candidate inpatient episodes. We excluded 1,222 episodes that lacked one or more key pathway
1460 nodes required for full-pathway assessment or had insufficient clinical documentation for reliable evaluation. This
1461 left 6,796 retained episodes spanning 287 disease entities. The retained cohort was used exclusively for retrospective
1462 multicentre validation.

1463 Because both cohorts were derived from authentic clinical records rather than heavily pre-curated benchmark cases,
1464 they retained substantial real-world heterogeneity and noise, including long-tail disease composition, uneven documen-
1465 tation quality and incomplete records at specific pathway stages. We deliberately preserved this complexity to better
1466 approximate real orthopaedic practice. As in other real-world clinical benchmarks, such heterogeneity substantially
1467 increases the difficulty of evaluation relative to narrowly filtered datasets.


1468 E.2 Prospective Dataset Curation

1469 In this section, we describe how the prospective cohorts were screened and allocated for the crossover reader study
1470 and the randomized clinical deployment study. The full workflow is summarized in Fig. E4. For both studies, cohort
1471 definition was performed at the patient or indexed-episode level, and only cases with complete longitudinal coverage
1472 of the predefined orthopaedic care pathway were retained.

1473 We first curated the prospective crossover reader-study cohort. Consecutive orthopaedic cases evaluated at Rui-
1474 jin Hospital between September 2025 and April 2026 were screened. Among 2,684 potentially eligible cases, inclusion
1475 required age ≥ 18 years, evaluation within the musculoskeletal care pathway, a complex musculoskeletal condition
1476 requiring sequential decision-making from admission to operation, available multimodal clinical data, and complete
1477 coverage of all four predefined study tasks: admission diagnosis, preoperative diagnosis, surgical planning and intra-
1478 operative diagnosis. We excluded 814 cases, including 123 with age < 18 years, 175 with insufficient imaging quality
1479 or missing key imaging studies, 116 with incomplete clinical documentation, 87 with missing operative records or
1480 intraoperative findings, 240 that did not complete all four study stages within the same episode of care, and 73 with-
1481 out definitive reference diagnosis or adequate follow-up. This yielded a final eligible full-chain cohort of 1,870 cases,
1482 all of which contained complete information for P1–P4. Reference labels for diagnostic tasks were derived from the
1483 final clinical record. Surgical-planning quality was not treated as a discrete reference label and was assessed sepa-
1484 rately by blinded senior orthopaedic surgeons using predefined criteria. Additionally, thirty orthopaedic physicians
1485 were prospectively recruited, stratified by experience level, and completed randomized crossover evaluation with and
1486 without OrthoPilot after a washout period of at least 6 weeks.

1487 We next curated the prospective randomized clinical deployment cohort. Consecutive orthopaedic inpatients admit-
1488 ted to Ruijin Hospital between September 2025 and April 2026 were screened. Among 9,742 potentially eligible
1489 participants, inclusion required age ≥ 18 years, admission to the orthopaedic inpatient service for diagnostic assess-
1490 ment or surgical treatment, availability of core multimodal clinical data, and completion of all 11 predefined key stages

A Internal Ruijin cohort curation



B External multicentre validation curation




Fig. E3: Cohort curation for OrthoBench and multicentre retrospective validation. **a**, Internal Ruijin cohort curation. Earlier cases were reserved for historical retrieval ($n = 48,464$). The remaining records yielded 89,746 patients for benchmark construction, from which patient-level allocation produced OrthoBench ($n = 5,905$); the remainder was used for training-cohort construction. **b**, External multicentre validation curation. Among 8,018 candidate inpatient episodes from 60 hospitals, 1,222 were excluded because key pathway information or sufficient documentation was lacking, leaving 6,796 retained episodes spanning 287 disease entities.

of the orthopaedic care pathway during the indexed admission. We excluded 1,502 patients, including 322 emergency trauma cases requiring immediate surgery, 241 with hospital stay shorter than 24 hours, 371 with incomplete electronic medical records or missing key multimodal data, 148 consultation-only or non-orthopaedic primary admissions, 79 duplicate admissions within the same treatment episode, and 341 who did not complete all 11 predefined care stages. This left 8,240 eligible patients, all of whom completed the full 11-stage orthopaedic care pathway. Patients were randomized in a 1:1 ratio using computer-generated randomization stratified by major disease category, surgical complexity and admission type, yielding 4,120 patients in the standard-care group and 4,120 in the OrthoPilot-assisted group. The standard-care group followed the conventional workflow, whereas OrthoPilot was integrated into the hospital information system in the intervention group to provide structured decision support across all 11 stages. Outcomes were prospectively assessed across hospital-level efficiency, physician workload, nursing workflow and patient-reported experience, and all analyses followed prespecified definitions.


E.3 Synthetic Dataset Curation

The five-stage OrthoPilot training curriculum draws on a combination of publicly available medical corpora, institutionally curated orthopaedic records, and synthetically generated reasoning trajectories. In this section, we describe the construction of synthetic question-answering data and tool-augmented trajectory data used across the training stages, complementing the raw clinical data described above. A central challenge in training a clinical reasoning model is the scarcity of high-quality, task-specific question-answer pairs that reflect the sequential decision structure of real patient care. To address this, we constructed synthetic question-answering datasets through two complementary pipelines: (1) *task-conditioned Question Answering (QA) generation* from the Ruijin Hospital orthopaedic cohort and (2) *multi-turn trajectory synthesis* for tool-augmented training.

Task-conditioned QA generation (Stage I–III).

For Stage I (clinical domain adaptation), we aggregated 6,565,111 medical instruction-style examples from 16 publicly available corpora, including public, curated private and synthetic sources, spanning general medicine, traditional Chinese medicine, biomedical literature and clinical guidelines (Table E3). These data were reformatted into a unified instruction–response schema to support broad medical knowledge acquisition without orthopaedic specialisation.

A Prospective crossover reader-study cohort curation



B Clinical deployment cohort curation




Fig. E4: Prospective cohort curation for the crossover reader study and randomized clinical deployment study. **a**, Prospective crossover reader-study curation. Among 2,684 potentially eligible orthopaedic cases screened at Ruijin Hospital between September 2025 and April 2026, 814 were excluded because of age younger than 18 years, insufficient imaging or missing key imaging studies, incomplete clinical documentation, missing operative records or intraoperative findings, incomplete coverage of the four predefined study stages, or lack of definitive reference diagnosis or adequate follow-up. The final eligible full-chain cohort contained 1,870 cases with complete information for P1–P4. **b**, Prospective randomized clinical deployment cohort curation. Among 9,742 potentially eligible orthopaedic inpatients screened between September 2025 and April 2026, 1,502 were excluded, leaving 8,240 eligible patients who completed the 11-stage orthopaedic care pathway. Patients were randomized 1:1 into the standard-care group ($n = 4,120$) and the OrthoPilot-assisted group ($n = 4,120$).

For Stage II (orthopaedic reasoning alignment), we derived 1,133,968 task-specific QA instances from the longitudinal Ruijin Hospital cohort by programmatically extracting patient-level contexts and pairing them with clinician-verified reference outputs across all 11 pathway tasks. Each instance was constructed by assembling the stage-appropriate clinical evidence (chief complaint, imaging reports, laboratory results, operative records or discharge summaries) as the input prompt, with the corresponding clinical decision (diagnosis, order set, surgical plan or rehabilitation recommendation) as the target response. To enrich reasoning depth, we supplemented these with external chain-of-thought medical reasoning data, including teacher-generated rationales from HuatuoGPT-o1 that were retained only when the associated final answer passed automated correctness filtering. The combined Stage II mixture totalled 35,242,453 instruction instances.

For Stage III (reinforcement learning), we curated 49,275 high-value training examples emphasising cases that require deeper reasoning, ambiguity resolution and clinically consequential decision-making. Approximately one-third were drawn from multi-disciplinary scientific reasoning tasks and two-thirds from the 11 orthopaedic tasks (~3,000 per task), ensuring broad coverage of both domain-specific and cross-disciplinary reasoning demands.

Multi-turn trajectory synthesis (Stages IV–V).

For the tool-augmented training stages, we synthesised multi-turn reasoning trajectories in which the model must actively invoke tools from both evidence worlds to solve clinical tasks. Short-horizon trajectories were generated by presenting high-quality orthopaedic task instances to multiple frontier teacher models, each of which independently produced candidate multi-turn reasoning paths with interleaved tool calls and evidence returns. Candidate trajectories were filtered by final-answer correctness and ranked by an LLM-based judge evaluating three dimensions: answer quality, appropriateness of tool use, and diversity of problem-solving strategy. Full-pathway trajectories concatenated

the 11 orthopaedic tasks into temporally ordered clinical decision chains, with the complete evidence-grounded trajectory from each preceding task retained as running context. Strict rejection sampling was applied, retaining only pathways in which all task answers were correct. Both trajectory types serve as the supervised corpus for Stage IV (multi-horizon clinical trajectory distillation), while Stage V (full-pathway agentic reinforcement learning) reuses the same task distribution under GRPO-style policy optimisation. A complete inventory of all training datasets, organised by training stage, is provided in Table E3.

Table E3: Training datasets used in each stage of the CHEESE curriculum, with sample counts and source links. Institutionally derived datasets are sourced from the Ruijin Hospital orthopaedic cohort; synthetically generated trajectory datasets are noted accordingly.

Dataset Name	Samples	Link / Source
Stage I: Clinical Domain Adaptation (Continued Pre-Training, (CPT))		
AGCT [1] (medical text)	421,216	Synthetic Data
Apollo [2] (medical QA)	3,482,252	https://huggingface.co/datasets/FreedomIntelligence/ApolloCorpus
BioLeaflets [3] (biomedical text)	1,336	Synthetic Data
biology_textbook [4] (text)	1,083	Synthetic Data
biology_wiki [5] (text)	1,083	Synthetic Data
biorxiv-clustering-s2s [6] (biomedical)	53,787	https://huggingface.co/datasets/mteb/biorxiv-clustering-s2s
ChatMed-Consult-Dataset [7] (medical QA)	549,317	https://huggingface.co/datasets/michaelwzhu/ChatMed-Consult-Dataset
Clinical_Guidelines [8] (guidelines)	37,970	Curated from public clinical guidelines
private data (medical text)	86,929	Synthetic Data
gliner-bird-diet-synthetic [9] (named entity recognition)	1,892	https://huggingface.co/datasets/wjbmattngly/gliner-bird-diet-synthetic
med_qa [10] (medical QA)	1,066,824	https://huggingface.co/datasets/bigbio/med_qa
Medical_Cord19 [11] (biomedical text)	300,000	https://huggingface.co/datasets/allenai/cord19
medicine_wiki [12] (medical text)	1,086	Synthetic Data
medqa_corpus_en [13] (medical QA)	76,043	https://huggingface.co/datasets/cogbuji/medqa_corpus_en
ReMeDi [14] (medical dialogue)	371,728	https://huggingface.co/datasets/ljvmiranda921/remedial-xsd-sft
ShenNong-TCM-Dataset [15] (traditional chinese medicine)	112,565	https://huggingface.co/datasets/michaelwzhu/ShenNong-TCM-Dataset
<i>Stage I Total</i>		6,565,111
Stage II: Musculoskeletal Reasoning Cold-Start Alignment (Supervised Fine-Tuning, (SFT))		
Bone Task 1 (admission diagnosis)	157,899	Ruijin Hospital orthopaedic cohort
Bone Task 2 (preoperative diagnosis)	157,917	Ruijin Hospital orthopaedic cohort
Bone Task 3 (Intraoperative diagnosis)	154,702	Ruijin Hospital orthopaedic cohort
Bone Task 4 (discharge diagnosis)	157,572	Ruijin Hospital orthopaedic cohort
Bone Task 5 (perioperative assessment)	180,441	Ruijin Hospital orthopaedic cohort
Bone Task 6 (surgical planning)	71,085	Ruijin Hospital orthopaedic cohort
Bone Task 7 (preoperative orders)	52,336	Ruijin Hospital orthopaedic cohort
Bone Task 8 (postoperative orders)	120,949	Ruijin Hospital orthopaedic cohort
Bone Task 9 (discharge summary)	52,320	Ruijin Hospital orthopaedic cohort
Bone Task 10 (rehabilitation)	8,418	Ruijin Hospital orthopaedic cohort
Bone Task 11 (MDT consultation)	20,329	Ruijin Hospital orthopaedic cohort
External medical QA corpora	—	Multiple public sources (see Methods)
<i>Stage II Bone Subtotal</i>		1,133,968
<i>Stage II Total (incl. external)</i>		35,242,453
Stage III: Reason-Driven Reinforcement Learning (GRPO)		
Bone Tasks 1–11 (orthopaedic RL)	~33,000	Ruijin Hospital orthopaedic cohort
MegaScience [16] (multi-disciplinary reasoning)	~16,275	Curated scientific reasoning corpus
<i>Stage III Total</i>		49,275
Stage IV: Multi-Horizon Clinical Trajectory Distillation (SFT)		
Short-horizon orthopaedic agent trajectories	~82,000	Synthetic multi-turn task-level trajectories
Long-horizon full-pathway trajectories	~53,000	Synthetic full-pathway concatenated trajectories
Miroverse [17] (multi-disciplinary agent data)	~22,000	Curated agent-style corpora
<i>Stage IV Total</i>		~157,000 cases; 2.169B tokens

Dataset Name	Samples	Link / Source
Stage V: Full-Pathway Agentic Reinforcement Learning (GRPO)		
Bone Tasks 1–11 (full-pathway continuation queries)	~11,000	Ruijin Hospital orthopaedic cohort
<i>Stage V Total</i>	11,000	

Appendix F Implementation of CHEESE

F.1 Details of Backbone

CHEESE is built upon Qwen3-32B [18], a 32-billion-parameter dense Transformer language model from Alibaba’s Qwen series. Qwen3-32B employs 64 Transformer layers with a hidden dimension of 5,120 and 40 attention heads, using Grouped Query Attention with 8 key-value heads. The model supports a native context length of 32,768 tokens, which can be extended to 131,072 tokens via YaRN-based rotary position embedding interpolation. It uses SwiGLU activation functions in the feed-forward layers and RMSNorm for layer normalisation. The tokenizer is a byte-level byte pair encoding tokenizer with a vocabulary size of 151,936 tokens, providing comprehensive coverage of English, Chinese and multilingual medical terminology.

F.2 Implementation of Training Details

We describe the training configuration for each of the five stages in the CHEESE training curriculum. All experiments were conducted on NVIDIA A100 80 GB GPUs using distributed data parallelism, with mixed-precision training in bfloat16 throughout. A summary of all stage-level hyperparameters is provided in Table F4.

Stage I: Clinical domain adaptation.

Continuous pre-training was performed on the Qwen3-32B base model using the large-scale clinical instruction corpus described above. This stage was designed to adapt the general-purpose backbone to the medical domain and to establish broad clinical language understanding before task-specific specialization. Training employed Low-Rank Adaptation (LoRA) with rank $r = 16$ applied to all linear layers, a maximum sequence length of 4,096 tokens, a learning rate of 1×10^{-4} with cosine scheduling and 10% warmup, a per-device batch size of 1 with gradient accumulation over 8 steps, and training for 2 epochs.

Stage II: Musculoskeletal reasoning cold-start alignment.

Supervised fine-tuning was performed on the Stage I checkpoint using the orthopaedic reasoning and medical instruction data described in the preceding sections. This stage aimed to align the model with orthopaedic clinical workflows and to provide a cold-start policy capable of producing structured, task-specific reasoning across the full decision pathway. Training used LoRA with rank $r = 16$ on all layers, a maximum sequence length of 8,192 tokens, a learning rate of 1×10^{-4} with cosine scheduling and 10% warmup, a per-device batch size of 1 with 8-step gradient accumulation across 3 nodes, and training for 2 epochs.

Stage III: Reason-driven reinforcement learning with GRPO.

The training prompts were sampled from the curated reasoning data introduced above, covering both general scientific reasoning and orthopaedic clinical decision-making tasks. For each training prompt, $G = 8$ candidate responses were sampled from the current policy. Training used an actor learning rate of 5×10^{-6} , a train batch size of 1,024 prompts, a maximum prompt length of 1,024 tokens, a maximum response length of 4,096 tokens, a PPO mini-batch size of 128, tensor model parallelism with degree 2, and Kullback–Leibler Divergence (KL) regularisation with coefficient $\beta_{\text{KL}} = 0.001$ estimated using a low-variance KL estimator. Training was run for 5 epochs.

Stage IV: Multi-horizon clinical trajectory distillation.

Supervised fine-tuning on multi-turn, tool-augmented trajectories was performed at two complementary temporal scales. Short-horizon trajectories focused on individual musculoskeletal tasks and were selected according to final-answer correctness, answer quality, tool-use appropriateness and problem-solving diversity. Long-horizon trajectories connected all 11 tasks into temporally ordered full-pathway decision chains, enabling the model to learn longitudinal clinical reasoning across admission, diagnosis, treatment, postoperative care and rehabilitation planning. Training was conducted with full-parameter fine-tuning, a maximum context length of 32,768 tokens to accommodate multi-turn tool interactions, and standard supervised next-token prediction loss applied to reasoning traces and tool-call tokens while masking tool-returned observations.

Table F4: Summary of training hyperparameters across the five-stage CHEESE curriculum. Stages I–III operate in a single-turn, tool-free setting; Stages IV–V operate in a multi-turn, tool-augmented setting. LR: learning rate; BS: effective batch size; Seq. Len.: maximum sequence length.

Stage	Method	LR	BS	Seq. Len.	Epochs	Key details
I	LoRA CPT	1×10^{-4}	24	4,096	2	Rank 16, all layers
II	LoRA SFT	1×10^{-4}	24	8,192	2	Rank 16, all layers
III	GRPO RL	5×10^{-6}	1,024	4,096	5	$G=8$, $\beta_{\text{KL}}=0.001$
IV	Full SFT	1×10^{-5}	32	32,768	2	Short + long-horizon trajectories
V	GRPO RL	5×10^{-6}	1,024	32,768	5	Full-pathway agentic RL

Stage V: Full-pathway agentic reinforcement learning.

Building upon the trajectory representations distilled in Stage IV, CHEESE was further optimised as a clinical agent through trajectory-level reinforcement learning. The reward function combined three terms: $R(H_T) = \alpha R_{\text{acc}}(H_T) + \beta R_{\text{fmt}}(H_T) + \gamma R_{\text{tool}}(H_T)$, where R_{acc} is the task-accuracy reward, R_{fmt} is the format-compliance reward, and R_{tool} is the tool-use reward encouraging active and appropriate invocation of tools from both evidence worlds. Training used a mixed regime of single-task queries and full-pathway continuation queries, optimised under GRPO-style updates with the same distributed infrastructure as Stage III.

Appendix G Baselines

To rigorously evaluate OrthoPilot across the full orthopaedic care pathway, we benchmarked it against 20 baselines selected to represent the current frontier of general-purpose, medical-specific, and agentic artificial intelligence systems. These baselines were stratified into four categories: 6 proprietary frontier models, 9 open-source or openly released general-purpose models, 3 medical-specific models, and 2 agentic framework baselines.

G.1 Proprietary Models

Claude-Sonnet-4.5. Claude Sonnet 4.5 is a closed-source general-purpose model developed by Anthropic and officially released in September 2025 [19]. It was designed as a frontier production model with strong reasoning and agentic computer-use capabilities, making it an appropriate commercial baseline for comparison with a clinically specialized musculoskeletal system.

GPT-4o. GPT-4o is a closed-source multimodal model developed by OpenAI and released in May 2024 [20]. It has become a widely used industrial baseline in recent evaluations of general-purpose large language models. It therefore provides a strong proprietary comparator without task-specific optimization for musculoskeletal workflows.

GPT-5.1. GPT-5.1 is treated here as a closed-source OpenAI model version accessed through the OpenAI API [21]. It serves as a newer-generation proprietary baseline for assessing whether performance gains arise from frontier general-model capability alone rather than from musculoskeletal-specific system design.

Gemini-2.5-flash. Gemini 2.5 Flash is a closed-source model developed by Google and announced in April 2025 as a preview model emphasizing low latency, cost efficiency, and controllable reasoning [22]. It is a realistic commercial baseline for deployment-oriented settings because many practical clinical applications rely on fast general models rather than slower specialist systems.

Grok-4. Grok 4 is a proprietary general-purpose model developed by xAI [23]. It is a reasonable baseline because it broadens comparison across major commercial model providers and reduces dependence on any single API ecosystem.

Mistral-Large. Mistral Large is a closed-source flagship model developed by Mistral AI and released in February 2024 [24]. Although Mistral AI also provides open-weight models, the Mistral Large line itself is proprietary and therefore captures the performance of a leading commercial model family in a non-specialized setting.

G.2 Open-Source General-Purpose Models

This subsection includes open-source and openly released general-purpose models with publicly available technical reports or model artifacts.

DeepSeek-R1. DeepSeek-R1 is an openly released general-purpose reasoning model developed by DeepSeek-AI and released in January 2025 [25]. It builds on large-scale reinforcement learning to elicit extended reasoning behaviour, making it a strong open comparator for complex multi-step clinical tasks.

DeepSeek-V3. DeepSeek-V3 is a general-purpose mixture-of-experts language model developed by DeepSeek-AI and released in December 2024 [26]. With 671 billion total parameters and 37 billion activated parameters per token, it serves as a competitive, openly released non-medical baseline for broad instruction following, reasoning, and text generation.

1627 **GPT-oss-120B.** GPT-oss-120B is an openly released general-purpose language model developed by OpenAI and
1628 released in August 2025 [27]. It is a sparse mixture-of-experts model designed to provide strong reasoning and general
1629 instruction-following performance in an open deployment setting.

1630 **Llama-3.3-70B.** Llama-3.3-70B-Instruct is an openly available general-purpose instruction-tuned model from the
1631 Llama family [28]. It provides a strong, widely adopted community baseline for multilingual generation and reasoning
1632 and therefore reflects the performance ceiling achievable with a popular non-medical open-weight model at this scale.

1633 **Llama-4-Maverick.** Llama-4-Maverick is an openly available general-purpose mixture-of-experts model released as
1634 part of the Llama 4 family in 2025 [29]. It represents a newer-generation open-weight architecture for general reasoning
1635 and instruction following without specialty adaptation.

1636 **Phi-4.** Phi-4 is an openly released dense language model developed by Microsoft and released in December 2024 [30].
1637 Despite its relatively compact 14-billion-parameter scale, it was designed to deliver strong reasoning quality relative
1638 to model size, making it a useful efficiency-oriented baseline.

1639 **Qwen3-235B-A22B.** Qwen3-235B-A22B is an openly released general-purpose mixture-of-experts model developed
1640 by the Qwen team at Alibaba and released in April 2025 [31]. With 235 billion total parameters and 22 billion activated
1641 parameters, it represents a large contemporary open model with broad general capabilities but no musculoskeletal
1642 specialization.

1643 **Qwen3-32B.** Qwen3-32B is an openly released general-purpose dense model from the same Qwen3 family [31]. It offers
1644 a mid-scale counterpart to larger frontier open models while retaining strong general reasoning and instruction-following
1645 performance.

1646 **Kimi-K2.** Kimi-K2 is an open-source general-purpose mixture-of-experts model developed by Moonshot AI and
1647 released in July 2025 under a modified MIT license with publicly available code and model weights [32]. It is an infor-
1648 mative comparator because it represents a strong open agentic model with broad reasoning and tool-use capabilities
1649 while remaining a non-medical system without musculoskeletal specialization.

1650 G.3 Medical-Specific Models

1651 **HuatuoGPT-o1-72B.** HuatuoGPT-o1-72B is a medical-specific large language model developed to improve complex
1652 clinical reasoning rather than general-domain dialogue alone [33]. Introduced in December 2024, it combines verifiable
1653 medical problems, supervised reasoning traces, and reinforcement learning with verifier-based rewards to enhance
1654 multi-step medical decision-making.

1655 **MedGemma-27B.** MedGemma-27B is a medical foundation model released by Google in July 2025 as part of the
1656 MedGemma family [34]. Built on the Gemma 3 architecture, it was designed for medical understanding and reasoning
1657 across clinical text and images while preserving the broad capabilities of the underlying general model.

1658 **OpenBioLLM-70B.** OpenBioLLM-70B is an open biomedical language model developed by Saama AI Labs in 2024
1659 on top of Llama 3 70B [35]. It serves as a reasonable baseline because it represents a large-scale biomedical model
1660 intended for healthcare and life-science applications, thereby testing whether broad biomedical adaptation alone is
1661 sufficient for complex musculoskeletal tasks.

1662 G.4 Agentic Framework

1663 **MDAgents.** MDAgents is an agentic framework for medical decision-making introduced at NeurIPS 2024, which
1664 adaptively assigns solo or group collaboration structures to multiple large language model agents according to task
1665 complexity [36]. It is a relevant comparator because it tests whether structured multi-agent orchestration improves
1666 clinical reasoning beyond single-model inference.

1667 **MDAgents-GPT4o.** MDAgents-GPT4o denotes our instantiation of the same framework with GPT-4o as the back-
1668 bone model [20, 36]. This variant evaluates whether the framework can realize meaningful gains when paired with a
1669 strong proprietary general-purpose model.

1670 **MDAgents-DeepSeek.** MDAgents-DeepSeek denotes our instantiation of the same framework with a DeepSeek-
1671 V3 backbone [26, 36]. This variant assesses whether the framework remains competitive when instantiated with a
1672 high-performing, openly released model rather than a proprietary backbone.

1673 Appendix H OrthoBench Design and Evaluation Metrics




Fig. H5: Disease Distribution and Taxonomy in OrthoBench. Radial cladogram displaying all 1,000 benchmark disease codes organised by International Classification of Diseases, Tenth Revision (ICD-10) hierarchy: ICD chapter (innermost coloured arc) → ICD category (3-character code) → individual disease (leaf node). Each chapter is colour-coded according to its ICD-10 classification. Two concentric annotation bands surrounding the leaf nodes encode distributional status (blue = in-distribution (ID); pink = out-of-distribution (OOD)) and disease prevalence (purple = rare; grey = non-rare). Leaf node size is proportional to patient count. Disease names are displayed radially at the outermost ring. The dendrogram reveals that out-of-distribution and rare diseases are non-uniformly distributed across ICD-10 chapters, with malignant neoplasms (chapter C) and congenital malformations (chapter Q) harbouring the highest proportions of rare and OOD codes.




Fig. H6: Overview of OrthoBench Composition. **a**, Distribution of evaluation instances across the 11 clinical tasks, grouped by category (pink = diagnostic reasoning, blue = clinical management). The dashed line indicates the mean sample count across tasks. Diagnostic tasks each contain approximately 17,700 instances, whereas management tasks range from 5,563 to 20,464 instances. **b**, Generalisation strata defined by distributional status (ID versus OOD) and prevalence (non-rare versus rare), showing disease codes (blue bars, left axis) and patients (pink bars, right axis). The OOD×Rare stratum contains 200 codes and 360 patients. **c**, Distribution of the 1,000 benchmark disease codes across ICD-10 chapters. Chapters with fewer than 10 codes are grouped as “Other”. **d**, Long-tail frequency distribution of patient counts per disease code on a log scale, coloured by distributional status (blue = ID, pink = OOD). **e**, Number of disease codes per ICD-10 chapter stratified by distributional status (blue = ID, pink = OOD). **f**, Number of disease codes per ICD-10 chapter stratified by prevalence (blue = non-rare, pink = rare).

Table H5: Overview of OrthoBench Clinical Tasks. Each of the 11 clinical tasks is mapped to a category (diagnostic reasoning or clinical management), together with its clinical description, applicable question types (Open = open-ended generation; MC = four-option multiple-choice; TF = true-or-false verification), total number of evaluation samples, and the number of unique patients. Diagnostic tasks employ all three question formats, whereas management tasks are evaluated exclusively as open-ended generation. Total evaluation instances across all tasks sum to 135,745.

Task	Category	Clinical Description	Question Types	Samples	Patients
1	Diagnostic	Admission diagnosis	Open / MC / TF	17,713	5,905
2	Diagnostic	Preoperative diagnosis	Open / MC / TF	17,715	5,905
3	Diagnostic	Intraoperative diagnosis	Open / MC / TF	17,400	5,905
4	Diagnostic	Discharge diagnosis	Open / MC / TF	17,712	5,905
5	Management	Perioperative assessment	Open	20,464	5,905
6	Management	Surgical planning	Open	8,154	5,905
7	Management	Preoperative orders	Open	5,875	5,875
8	Management	Postoperative orders	Open	13,665	5,905
9	Management	Discharge summary	Open	5,866	5,866
10	Management	Rehabilitation planning	Open	5,618	5,618
11	Management	Multidisciplinary consultation	Open	5,563	5,563
Total				135,745	5,905

Table H6: Generalisation Strata of OrthoBench. Disease codes and test patients are cross-tabulated according to distributional status (ID versus OOD) and disease prevalence (non-rare versus rare), resulting in four evaluation strata that together comprise 1,000 disease codes and 5,905 patients.

	Non-rare		Rare		Total	
	Codes	Patients	Codes	Patients	Codes	Patients
ID	321	3,051	62	274	383	3,325
OOD	417	2,220	200	360	617	2,580
Total	738	5,271	262	634	1,000	5,905

H.1 Task Design

OrthoBench comprises 11 clinical tasks that collectively trace the full decision pathway of an orthopaedic inpatient encounter, from the moment a patient presents to final post-discharge management (Table H5). Rather than sampling questions from textbooks or licensing examinations, every evaluation instance in OrthoBench is derived from a real patient record within a longitudinal cohort of over 180,000 orthopaedic admissions spanning two decades (2004–2024) at a single tertiary centre. This design ensures that each question is anchored to authentic clinical complexity, including comorbidity profiles, imaging findings, and laboratory trajectories that would be absent from synthetic or textbook-based benchmarks.

The 11 tasks are partitioned into two clinically motivated categories. The first category, diagnostic reasoning (Tasks 1–4), evaluates a model’s capacity to synthesise progressively richer clinical information into an accurate diagnosis. Task 1 (admission diagnosis) tests inference from the chief complaint, history of present illness, and basic demographics alone. Task 2 (preoperative diagnosis) introduces imaging reports, laboratory results, and physical examination findings, requiring the model to refine or revise its initial impression. Task 3 (Intraoperative diagnosis) provides intraoperative findings, such as arthroscopic observations or histopathological frozen-section results, that may alter the working diagnosis. Task 4 (discharge diagnosis) presents the complete hospital course and expects a final, consolidated diagnostic formulation. Each of these four tasks is administered in three question formats: open-ended generation, four-option multiple-choice, and true-or-false verification, yielding approximately 17,700 evaluation instances per task and enabling assessment of both recall and discrimination facets of diagnostic competence.

Table H7: Disease distribution by ICD-10 chapter. Disease codes and test patients are stratified by ICD-10 chapter, distributional status (ID versus OOD), and disease prevalence (non-rare versus rare).

Ch.	Chapter Name	Codes	Patients	ID	OOD	Non-rare	Rare
M	Musculoskeletal & connective tissue	309	1,949	127	182	260	49
S	Injuries (single body region)	254	2,290	109	145	245	9
D	Neoplasms & blood disorders	115	216	34	81	55	60
C	Malignant neoplasms	106	426	39	67	27	79
T	Injuries (multiple) & complications	50	102	19	31	47	3
Q	Congenital malformations	47	123	16	31	11	36
R	Symptoms & signs	26	80	9	17	26	0
G	Nervous system diseases	21	60	7	14	17	4
L	Skin & subcutaneous tissue	17	20	2	15	14	3
A	Infectious diseases	14	17	6	8	3	11
Z	Health services contact	13	590	9	4	13	0
E	Endocrine & metabolic	11	14	4	7	6	5
I	Circulatory system	10	10	2	8	7	3
J	Respiratory system	3	4	0	3	3	0
K	Digestive system	3	3	0	3	3	0
F	Mental & behavioural	1	1	0	1	1	0
Total		1,000	5,905	383	617	738	262

The second category, clinical management (Tasks 5–11), shifts the evaluation target from diagnostic accuracy to the generation of executable, patient-specific clinical plans. Task 5 (perioperative assessment) requires identification and quantification of surgical risk factors, such as cardiovascular comorbidities, anticoagulant use, diabetes control, and anaesthesia considerations, producing 20,464 evaluation instances across the full test cohort. Task 6 (surgical planning) tests the ability to recommend an appropriate operative strategy, including implant choice and surgical access, given the confirmed diagnosis and patient-specific anatomy. Task 7 (preoperative orders) assesses the generation of structured preoperative management plans encompassing surgical preparation, medication adjustment, fasting instructions, perioperative evaluation, and required examinations. Task 8 (postoperative orders) assesses the generation of immediate postoperative care plans encompassing analgesia regimens, thromboprophylaxis, antibiotic protocols, mobilisation instructions, and monitoring parameters. Task 9 (discharge summary) requires synthesis of the entire hospital course into a concise clinical narrative suitable for the referring physician. Task 10 (rehabilitation planning) evaluates the formulation of stage-appropriate physiotherapy and functional recovery programmes tailored to the specific procedure and patient profile. Task 11 (multidisciplinary consultation) tests the model’s ability to identify scenarios requiring specialist input, such as endocrinology review for metabolic bone disease or cardiology clearance for high-risk patients, and to draft an appropriately directed consultation request.

All management tasks are evaluated exclusively in an open-ended generation format, reflecting the clinical reality that real-world care plans are free-form documents rather than fixed-choice selections. The total benchmark comprises 135,745 evaluation instances across 5,905 unique test patients (Table H5, Fig. H6a). Task sample counts vary from 5,563 (Task 11, multidisciplinary consultation) to 20,464 (Task 5, perioperative assessment), reflecting natural variation in the availability of structured ground-truth labels across different clinical documentation types.

H.2 Data Composition

OrthoBench includes 1,000 distinct disease codes drawn from the ICD-10, covering a broad spectrum of musculoskeletal, trauma, oncological and systemic conditions encountered in orthopaedic care. Codes were organised by ICD-10 chapter, three-character category and individual disease code (Table H7, Fig. H5).

Disease codes were concentrated in musculoskeletal and connective tissue diseases (chapter M; 309 codes, 1,949 patients) and injuries to a single body region (chapter S; 254 codes, 2,290 patients), which together accounted for 56.3% of benchmark disease codes and 71.8% of test patients (Table H7, Fig. H6c). The next largest chapters were neoplasms and blood disorders (chapter D; 115 codes) and malignant neoplasms (chapter C; 106 codes). Smaller chapters included trauma complications (chapter T; 50 codes), congenital malformations (chapter Q; 47 codes), and disorders of the nervous system (chapter G; 21 codes). The remaining nine chapters collectively contributed 98 codes.

For generalisation analysis, disease codes were stratified by distributional status and prevalence, yielding 383 in-distribution and 617 out-of-distribution codes, and 738 non-rare and 262 rare codes (Table H6, Fig. H6b,e,f).

All 5,905 test patients contributed to the four diagnostic tasks and to Tasks 5, 6 and 8. Tasks requiring specialised documentation had smaller sample sizes: 5,875 patients for preoperative orders, 5,866 for discharge summary generation, 5,618 for rehabilitation planning and 5,563 for multidisciplinary consultation (Table H5).

Algorithm 2 ORACLE scoring for open-ended clinical management tasks

```
1: Input: model answer  $A$ ; reference answer  $G$ ; physician rubric  $B$ ; judge set  $\mathcal{J}$ ; level weights  $w_l, l \in \mathcal{L}$ 
2: Output: normalized ORACLE score  $s$  and aggregated summaries
3:  $\mathcal{L} \leftarrow \{\text{primary, secondary, additional}\}$ 
4:  $\mathcal{C} \leftarrow \text{Scorecard}(G, B)$ , organised as category-level item lists
5:  $\mathcal{E} \leftarrow \{e_{ilj} : e_{ilj} \in \mathcal{C}_{i,l}, e_{ilj} \neq \emptyset, l \in \mathcal{L}\}$ 
6: for  $e_{ilj} \in \mathcal{E}$  do
7:   for  $J_k \in \mathcal{J}$  do
8:      $z_{ilj}^{(k)} \leftarrow J_k(A, e_{ilj}, B_i) \in \{0, 1\}$ 
9:   end for
10:   $c_{ilj} \leftarrow \mathbf{1} \left[ \sum_{k=1}^{|\mathcal{J}|} z_{ilj}^{(k)} \geq \lfloor |\mathcal{J}|/2 \rfloor + 1 \right]$ 
11: end for
12:  $S_{\text{num}} \leftarrow \sum_i \sum_{l \in \mathcal{L}} \sum_{j=1}^{n_{il}} w_l c_{ilj}$ 
13:  $S_{\text{den}} \leftarrow \sum_i \sum_{l \in \mathcal{L}} n_{il} w_l$ 
14:  $s \leftarrow S_{\text{num}}/S_{\text{den}}$ 
15: Report  $s$  with case-level, patient-level, category-level and level-specific summaries
16: return  $s$ 
```

Table H8: Demographics of ORACLE physician evaluators. Demographic and professional characteristics of the 12 orthopaedic specialists who participated in the evaluation. All physicians were from tertiary (Grade A) hospitals with diverse subspecialty expertise, ensuring comprehensive clinical perspective across the evaluation tasks.

Characteristic	Count	Percentage
Professional Title		
Resident Physician	2	16.7%
Attending Physician	4	33.3%
Associate Chief Physician	3	25.0%
Chief Physician	3	25.0%
Years of Experience		
<5 years	2	16.7%
5–10 years	4	33.3%
10–20 years	4	33.3%
>20 years	2	16.7%
Subspecialty		
General Orthopaedics	3	25.0%
Trauma Orthopaedics	3	25.0%
Spine Surgery	2	16.7%
Joint Surgery	2	16.7%
Rehabilitation Medicine	2	16.7%
Hospital Level		
Tertiary Grade A Hospital	12	100.0%

H.3 Physician Evaluation of ORACLE

Evaluation design and physician cohort.

We collected 2,073 blinded evaluations across 691 clinical cases, with three independent physician reviews per case (evaluator demographics are summarised in Table H8). The cohort covered the seven open-ended management tasks and included four candidate systems: CHEESE, GPT-5.1, MedGemma-27B, and DeepSeek-R1. Each physician ranked the four model responses from best to worst, scored them on five Likert dimensions, and marked critical safety errors when present (Fig. H7). Evaluation time was also recorded for each review.

Agreement between ORACLE and physician rankings.

Physician rankings showed strong alignment with ORACLE scores. Across 8,292 model-response evaluations, the overall Spearman correlation between automatic scores and physician ranks was -0.895 and the Pearson correlation




Fig. H8: Summary of physician evaluation results for ORACLE. **a**, Task-specific Bradley–Terry preference parameters across the seven management tasks and overall performance. CHEESE shows the highest overall preference parameter, whereas MedGemma-27B leads surgical planning and GPT-5.1 narrowly leads discharge summary generation. **b**, Mean Likert scores across accuracy, completeness, safety, actionability and clarity. **c**, Distribution of physician rankings across models. **d**, Critical error rates across models. N = 2,073 physician evaluations across 691 clinical cases.

Multi-dimensional quality ratings.

Multidimensional quality scoring showed a similarly structured pattern (Fig. H8b). CHEESE achieved the highest scores in accuracy (4.75), completeness (4.74), actionability (4.74), and clarity (4.71). GPT-5.1 performed best in safety and remained strong in clarity. MedGemma-27B showed a selective strength in actionability (4.68) despite lower overall scores. DeepSeek-R1 scored lowest across most dimensions. These results suggest that the physician signal captured by ORACLE was supported by broader quality judgments rather than by rank order alone.

Ranking consistency across cases.

The ranking distribution reinforced these differences (Fig. H8c). CHEESE was ranked first in 48.6% of evaluations and had the best mean rank of 1.80. GPT-5.1 ranked first in 27.7% of evaluations and had a mean rank of 2.23. MedGemma-27B ranked first in 15.4% of evaluations and was concentrated in third place. DeepSeek-R1 ranked fourth in 60.2% of evaluations and had the poorest mean rank of 3.32. Together, these distributions indicate that CHEESE was not only strong on average but also more consistently preferred across cases.

Inter-rater concordance and evaluation time.

Concordance among physicians was quantified with Kendall’s coefficient of concordance (W), the recommended statistic for ordinal ranking data with multiple raters. Overall W reached 0.86 and per-task values ranged from 0.83 to 0.89, indicating strong concordance across the seven management tasks (Table H9). Mean evaluation time was 142.1 s and varied narrowly between 123.9 s and 152.1 s. Three observations follow. The high concordance values indicate that physician rankings reflect a shared latent quality signal rather than idiosyncratic preference, supporting the use of these rankings as ground truth for ORACLE’s correlation analysis. The narrow range of review times indicates that no task imposed disproportionate cognitive load, which supports the scalability of the protocol to larger benchmarks. Together, these findings establish that the physician evaluation panel produced consistent and comparable judgments across tasks, providing a stable reference against which the automated ORACLE pipeline can be calibrated.

Table H9: Inter-rater concordance and evaluation time across clinical tasks. Kendall’s coefficient of concordance (W) was computed across physician rankings within each task. Mean evaluation time reflects the cognitive complexity of each clinical assessment.

Task Name	Kendall’s W	Mean Time (s)	n Cases
Perioperative Assessment	0.89	123.9	100
Preoperative Orders	0.87	142.4	100
Surgical Planning	0.84	152.1	100
Postoperative Orders	0.86	147.2	100
Discharge Summary	0.87	143.4	100
Rehabilitation Planning	0.83	148.8	91
Multidisciplinary Consultation	0.86	137.9	100
Overall	0.86	142.1	691

1774 Critical safety errors.

1775 To assess clinical safety beyond subjective ratings, we quantified critical errors, defined as model outputs containing
1776 recommendations that could directly harm patients if implemented without physician oversight (Fig. H8d). Physician
1777 evaluators identified errors including contraindicated medications, inappropriate surgical approaches, missed critical
1778 diagnoses and unsafe discharge recommendations. GPT-5.1 showed the lowest critical error rate at 1.0%. CHEESE
1779 followed closely at 1.2%, demonstrating that multi-stage training of a 32B parameter domain-specific model can
1780 effectively increase orthopaedic knowledge density and reduce error rates to levels comparable with generalist frontier
1781 systems. MedGemma-27B showed a higher rate at 3.5%. DeepSeek-R1 exhibited the highest critical error burden at
1782 12.0%. These findings suggest that domain-specific optimisation through staged training can achieve safety profiles
1783 approaching state-of-the-art generalist models, while maintaining the clinical depth required for specialist tasks.

1784 Appendix I Details of Reader Study

1785 The retrospective reader study included 307 surgical orthopaedic patients sampled from the Ruijin Hospital longitu-
1786 dinal cohort from 2018 to 2023. These cases were excluded from OrthoBench and from CHEESE training. Sampling
1787 covered seven disease categories: trauma, degenerative joint disease, spine disorders, sports injuries, tumours, congen-
1788 ital and developmental conditions, and other orthopaedic diseases. Eighty-one orthopaedic physicians completed the
1789 study, with 27 Experts (≥ 25 years), 27 Specialists (10–24 years) and 27 Residents (1–9 years). In Phase 1, each physi-
1790 cian reviewed all 307 patients without AI support. After a one-month washout, the same physicians reviewed the same
1791 cases with OrthoPilot output available as a reference. OrthoPilot was also evaluated as a standalone system. Each case
1792 was assessed across eleven care-pathway tasks. T1–T4 were scored by accuracy, whereas T5–T11 were scored using
1793 ORACLE. The design produced 550,451 individual evaluations. Table I10 provides the complete table version of the
1794 retrospective human–AI comparative reader study. It reports the per-reader task scores, group means and OrthoPilot
1795 standalone results across all eleven care-pathway tasks.

Table I10: Individual physician performance in the retrospective human–AI comparative reader study. Values are mean scores (%) for each task. T1–T4 are closed-ended diagnostic tasks scored by accuracy. T5–T11 are open-ended management tasks scored by ORACLE. Overall denotes the mean across all eleven tasks. Reader IDs are anonymised physician identifiers with years of experience. The final row within each physician tier reports the mean.

Group	Reader	n	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	T11	Overall
OrthoPilot	AI	307	75.9	69.4	82.5	77.8	78.4	70.7	82.4	83.3	88.7	90.4	80.9	80.0
Expert	E18 (30y)	307	71.0	65.5	78.9	78.8	65.6	56.7	71.2	72.3	76.9	78.7	69.5	71.4
	E16 (35y)	307	79.6	64.0	79.3	62.4	65.7	58.2	71.7	71.9	75.9	78.2	67.1	70.4
	E07 (30y)	307	67.0	73.8	77.4	67.7	64.9	55.6	71.3	70.4	75.6	78.7	67.3	70.0
	E19 (35y)	307	60.1	60.3	71.9	58.4	63.6	58.4	71.4	72.4	76.6	80.7	67.2	67.4
	E27 (35y)	307	57.1	59.1	64.0	75.1	64.9	56.8	70.5	70.3	74.8	77.0	68.3	67.1
	E24 (33y)	307	72.2	60.9	54.9	64.4	66.0	55.6	70.8	70.1	75.4	75.5	67.8	66.7
	E13 (33y)	307	66.0	50.0	63.8	66.9	63.0	58.0	71.7	71.7	76.6	76.7	68.7	66.6
	E01 (30y)	307	67.2	60.0	57.9	64.1	65.6	56.9	68.5	70.2	75.2	75.3	68.2	66.3
	E14 (35y)	307	65.4	60.4	60.1	56.7	63.2	57.9	71.6	72.1	76.8	76.4	68.7	66.3
	E10 (25y)	307	76.4	63.1	57.1	56.1	65.0	57.0	68.9	68.8	74.3	74.9	66.2	66.2
	E03 (33y)	307	63.0	56.9	59.2	63.9	64.0	56.7	69.8	70.5	75.3	75.6	69.7	65.9
	E06 (30y)	307	62.8	60.6	65.2	60.0	62.9	55.3	68.3	68.2	75.2	76.5	67.6	65.7
	E12 (29y)	307	71.4	58.3	60.3	54.7	65.9	55.3	70.0	69.8	73.4	76.1	67.4	65.7
	E21 (30y)	307	71.6	59.4	71.3	45.0	64.1	57.7	67.7	69.9	75.8	75.0	65.5	65.7

(continued)

Group	Reader	n	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	T11	Overall
	E15 (30y)	307	59.7	50.1	70.3	63.3	63.0	55.0	70.5	69.8	75.3	77.3	66.2	65.5
	E22 (30y)	307	62.1	50.6	81.0	58.6	64.7	53.1	67.7	69.4	73.3	74.9	64.9	65.5
	E02 (35y)	307	68.5	62.4	64.1	48.9	63.3	56.8	68.4	69.8	74.7	76.2	65.3	65.3
	E04 (33y)	307	53.7	63.5	66.5	60.0	63.9	55.8	69.1	69.2	74.1	74.7	67.6	65.3
	E08 (26y)	307	61.8	58.4	56.0	64.0	63.2	55.9	68.7	71.3	72.9	75.5	66.8	65.0
	E11 (33y)	307	62.0	49.4	57.6	56.4	65.1	55.9	70.8	71.3	76.6	77.6	67.7	64.6
	E17 (36y)	307	59.9	59.9	46.8	66.3	63.0	53.0	69.7	71.1	75.0	76.2	67.5	64.4
	E09 (34y)	307	67.3	40.3	64.4	66.9	62.8	55.1	66.7	69.9	71.8	77.4	64.5	64.3
	E20 (25y)	307	72.7	43.5	67.9	52.3	63.6	54.8	66.7	70.2	73.7	73.7	65.4	64.0
	E23 (25y)	307	64.4	53.9	55.3	61.8	62.1	53.3	65.2	68.1	71.1	76.1	66.1	63.4
	E05 (30y)	307	57.0	50.4	48.7	45.2	62.2	53.2	67.4	69.0	71.4	74.7	63.8	60.3
	E26 (25y)	307	52.9	40.3	48.3	50.3	59.4	54.0	67.7	69.1	72.7	74.7	64.5	59.4
	E25 (29y)	307	56.8	39.8	44.8	54.6	59.4	51.8	66.4	65.9	71.0	73.1	64.0	58.9
	Mean	307	64.8	56.1	62.7	60.1	63.7	55.7	69.2	70.1	74.5	76.2	66.8	65.4
AI+Expert	E15 (30y)	307	65.9	67.6	80.7	76.7	84.3	76.1	87.2	86.3	95.4	93.2	97.4	82.8
	E07 (30y)	307	67.9	82.8	93.1	73.6	82.8	70.0	83.6	80.8	93.5	91.1	82.0	81.9
	E06 (30y)	307	67.5	68.2	73.3	77.1	80.6	73.4	86.6	90.8	96.4	94.6	90.2	81.7
	E13 (33y)	307	75.1	56.7	72.5	73.7	81.6	73.9	94.0	87.0	94.7	95.7	92.4	81.6
	E19 (35y)	307	61.1	62.2	81.0	59.3	84.8	83.2	85.8	90.9	94.4	94.5	91.4	80.8
	E14 (35y)	307	66.5	67.9	65.6	66.6	80.8	81.4	87.2	88.9	94.2	90.6	82.0	79.2
	E01 (30y)	307	72.4	68.5	67.4	69.6	82.8	76.0	83.4	81.8	88.9	93.9	85.6	79.1
	E20 (25y)	307	82.5	53.3	71.0	58.4	83.4	71.8	85.8	88.1	94.1	95.7	80.9	78.6
	E17 (36y)	307	76.6	67.6	60.0	85.7	77.6	67.0	80.8	83.2	92.7	88.6	84.0	78.5
	E10 (25y)	307	77.1	71.3	62.9	61.6	81.3	72.1	89.2	85.0	90.3	94.1	76.3	78.3
	E21 (30y)	307	86.8	61.0	70.1	65.3	72.0	74.1	82.9	83.3	85.3	93.3	85.8	78.2
	E12 (29y)	307	81.1	62.1	63.2	71.0	78.6	69.0	82.4	84.5	82.3	91.5	80.3	76.9
	E11 (33y)	307	68.9	66.8	64.3	66.8	82.8	71.1	84.6	83.5	85.0	86.0	81.5	76.5
	E16 (35y)	307	77.8	78.3	87.7	69.2	71.0	60.9	78.8	77.4	82.8	82.1	74.1	76.4
	E27 (35y)	307	63.4	69.8	67.4	78.3	77.5	68.3	79.8	83.1	89.1	87.0	74.0	76.2
	E24 (33y)	307	72.8	61.1	68.3	71.1	75.2	71.5	82.0	82.7	81.4	90.0	77.0	75.7
	E09 (34y)	307	68.8	49.9	71.1	65.1	78.8	72.5	87.4	82.5	87.0	92.3	76.0	75.6
	E03 (33y)	307	68.6	59.1	63.4	68.1	76.5	68.5	75.8	79.8	85.0	93.7	80.4	74.4
	E18 (30y)	307	71.3	72.7	78.1	79.7	66.9	61.2	73.1	72.9	77.4	79.3	71.3	73.1
	E23 (25y)	307	65.0	59.9	75.6	65.6	76.8	61.7	76.0	75.9	81.0	86.7	77.8	72.9
	E26 (25y)	307	60.3	49.5	69.2	58.4	78.4	71.1	83.6	79.8	89.1	83.7	78.0	72.8
	E22 (30y)	307	64.7	59.5	88.4	66.0	67.3	55.3	71.8	81.1	81.7	82.9	68.5	71.6
	E02 (35y)	307	72.7	62.7	74.4	51.7	73.1	61.0	81.8	74.9	82.5	82.1	70.4	71.6
	E04 (33y)	307	57.6	70.4	70.5	62.9	71.5	60.4	74.5	75.7	89.5	79.0	72.3	71.3
	E08 (26y)	307	66.5	63.1	65.9	68.9	70.4	58.2	69.4	78.5	85.7	82.1	74.0	71.2
	E25 (29y)	307	64.3	49.2	50.9	62.6	73.8	62.2	79.8	71.6	86.1	88.8	75.3	69.5
	E05 (30y)	307	58.1	67.7	56.2	46.5	65.5	62.8	70.4	71.4	78.4	85.1	68.3	66.4
	Mean	307	69.7	64.0	70.8	67.4	76.9	68.7	81.4	81.5	87.6	88.8	79.5	76.0
Specialist	S21 (20y)	307	65.7	66.9	69.5	66.0	61.5	63.3	63.0	68.5	70.2	73.0	65.0	66.6
	S20 (19y)	307	59.2	56.0	62.2	58.5	57.6	59.7	63.8	65.6	68.9	70.2	62.9	62.2
	S27 (16y)	307	57.0	55.5	58.7	56.3	59.9	57.3	63.3	65.2	68.0	70.5	64.2	61.4
	S07 (13y)	307	55.9	52.2	57.8	54.0	60.1	60.3	67.0	64.4	69.7	72.6	60.8	61.3
	S12 (11y)	307	59.0	55.5	61.1	58.0	59.6	54.9	63.8	61.4	68.7	70.3	60.9	61.2
	S01 (20y)	307	56.4	54.3	60.8	56.6	59.5	57.2	60.9	65.3	66.6	69.7	58.3	60.5
	S04 (23y)	307	56.8	55.7	60.9	55.9	56.7	59.5	60.4	61.6	66.7	68.0	59.9	60.2
	S13 (16y)	307	59.8	57.5	61.7	59.7	56.4	54.6	60.9	62.7	61.6	66.7	57.0	59.9
	S14 (13y)	307	55.6	53.8	58.5	55.1	57.9	58.4	61.8	62.6	64.6	69.9	60.8	59.9
	S15 (18y)	307	58.2	56.4	57.7	56.7	54.7	56.9	62.1	60.2	66.5	66.7	63.0	59.9
	S23 (24y)	307	59.0	57.3	59.7	56.6	54.9	54.4	59.5	62.2	65.8	68.3	59.8	59.8
	S24 (20y)	307	53.5	53.5	58.5	55.3	58.9	55.5	61.5	64.1	65.7	69.3	61.9	59.8
	S09 (24y)	307	54.1	53.0	57.7	55.7	57.0	56.2	60.4	63.9	68.4	67.5	61.5	59.6
	S05 (18y)	307	54.4	53.0	59.1	55.2	57.9	56.8	61.5	65.5	64.0	67.5	59.7	59.5
	S06 (10y)	307	57.5	55.1	60.3	56.4	55.6	56.2	57.7	61.2	65.3	67.5	60.0	59.3
	S18 (24y)	307	55.2	53.8	59.8	55.6	56.9	54.5	61.6	63.8	62.9	66.1	61.6	59.3
	S03 (12y)	307	57.3	55.4	60.8	55.2	59.2	53.9	56.2	63.3	64.7	66.5	59.7	59.3
	S19 (12y)	307	56.1	54.1	57.9	55.1	55.6	55.6	60.8	60.5	65.3	66.9	63.6	59.2
	S16 (22y)	307	53.2	53.0	58.4	53.2	58.3	55.8	59.4	63.8	65.7	67.2	59.8	58.9
	S26 (14y)	307	54.9	53.0	57.3	53.5	54.4	56.1	58.1	63.8	66.0	67.5	60.1	58.6
	S10 (22y)	307	52.8	52.9	55.6	52.8	54.4	55.4	62.1	62.9	65.0	64.8	63.2	58.4
	S25 (17y)	307	52.6	54.0	56.8	54.8	57.3	54.1	58.4	62.9	64.5	66.2	59.8	58.3
	S22 (15y)	307	53.3	51.3	57.0	52.8	56.3	56.2	59.1	61.0	64.4	67.3	58.7	57.9
	S17 (19y)	307	51.1	49.0	53.3	50.1	56.2	55.9	58.3	62.5	63.7	70.7	61.3	57.5
	S08 (12y)	307	52.4	50.1	53.3	49.7	57.7	55.0	61.4	62.4	64.4	66.5	59.0	57.4
	S02 (11y)	307	51.4	50.0	53.6	51.0	55.1	56.0	55.3	61.3	66.1	65.8	58.5	56.7
	S11 (14y)	307	43.1	43.0	46.0	45.1	54.7	53.0	54.9	61.1	63.2	68.2	58.0	53.7

(continued)

Group	Reader	n	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	T11	Overall
	Mean	307	55.4	53.9	58.3	55.0	57.2	56.4	60.5	63.1	65.8	68.2	60.7	59.5
AI+Specialist	S13 (16y)	307	77.9	66.0	60.9	77.1	80.8	74.1	85.4	87.4	88.2	94.1	85.6	79.8
	S12 (11y)	307	73.4	68.5	80.8	63.6	79.7	72.6	84.8	83.9	94.5	88.2	79.0	79.0
	S18 (24y)	307	64.4	65.9	81.3	69.4	78.5	68.2	84.3	87.5	88.0	90.4	85.1	78.5
	S17 (19y)	307	58.5	61.6	62.5	63.2	81.2	79.9	87.4	87.8	95.2	97.5	83.0	78.0
	S10 (22y)	307	62.2	66.2	77.2	61.8	76.8	70.8	83.5	85.0	93.8	91.7	85.7	77.7
	S24 (20y)	307	64.7	62.4	68.5	62.4	82.1	72.6	86.2	84.0	92.6	92.2	83.3	77.4
	S20 (19y)	307	67.5	67.6	75.0	77.8	73.2	70.6	78.4	83.6	85.9	91.6	79.8	77.4
	S27 (16y)	307	71.8	75.7	68.5	64.2	77.5	67.9	83.8	80.4	87.1	83.6	83.6	76.7
	S08 (12y)	307	68.8	58.6	58.0	63.7	80.6	73.1	85.4	87.1	91.2	94.5	81.9	76.6
	S21 (20y)	307	72.5	66.0	77.7	76.4	76.0	67.6	79.0	67.2	85.4	87.5	80.8	76.0
	S23 (24y)	307	68.5	64.0	68.5	74.7	74.3	64.9	83.6	85.0	84.4	86.6	80.0	75.9
	S25 (17y)	307	56.3	59.5	64.6	55.5	75.0	73.0	83.5	88.7	96.2	94.3	83.0	75.4
	S03 (12y)	307	64.8	67.8	79.5	73.6	75.2	65.7	76.4	82.7	84.4	84.9	72.8	75.3
	S05 (18y)	307	80.6	64.0	68.8	65.9	72.5	69.0	77.0	80.3	85.4	85.7	75.1	74.9
	S26 (14y)	307	61.8	64.6	69.2	60.3	73.5	67.0	78.8	83.9	86.8	89.4	77.6	73.9
	S09 (24y)	307	68.3	62.4	64.4	77.5	72.3	63.7	77.7	77.1	86.0	87.5	74.8	73.8
	S04 (23y)	307	69.6	63.4	72.3	66.9	68.1	71.7	78.0	80.0	65.7	89.9	78.3	73.1
	S06 (10y)	307	63.4	68.2	80.6	68.0	68.8	68.5	75.7	80.6	79.0	79.7	71.4	73.1
	S22 (15y)	307	63.2	53.4	66.7	62.7	72.3	69.3	81.6	78.6	84.3	95.3	76.4	73.1
	S15 (18y)	307	64.3	67.0	66.4	62.0	73.5	68.0	75.7	78.4	83.7	84.4	75.4	72.6
	S14 (13y)	307	65.9	59.8	61.3	60.5	76.4	67.4	78.4	80.0	83.2	87.1	75.1	72.3
	S01 (20y)	307	60.2	63.7	71.1	62.6	66.4	68.8	77.0	82.7	84.9	78.3	74.9	71.9
	S07 (13y)	307	61.9	60.6	61.5	60.3	74.4	66.4	81.6	83.9	88.7	70.8	78.3	71.7
	S19 (12y)	307	68.0	61.5	70.1	62.6	71.4	63.0	76.1	78.4	74.6	86.1	71.9	71.2
	S16 (22y)	307	62.3	62.5	67.6	60.0	68.1	68.8	75.9	81.2	78.4	80.9	71.7	70.7
	S02 (11y)	307	61.6	60.0	61.4	66.5	69.4	63.5	75.8	76.1	81.3	81.0	70.6	69.7
	S11 (14y)	307	59.6	52.2	60.8	52.1	68.1	61.3	74.4	78.1	81.8	79.6	73.6	67.4
	Mean	307	66.0	63.4	69.1	65.6	74.3	68.8	80.2	81.8	85.6	87.1	78.1	74.6
Resident	R05 (8y)	307	68.7	63.9	67.5	65.8	61.6	53.9	61.9	67.1	68.5	65.9	59.7	64.0
	R22 (3y)	307	65.4	63.8	68.2	64.3	60.4	53.2	63.5	64.4	64.4	62.7	66.4	63.3
	R23 (3y)	307	56.1	53.1	56.4	54.6	63.5	50.4	59.8	63.9	66.3	63.5	62.9	59.1
	R09 (7y)	307	54.2	51.5	56.4	54.1	62.8	53.7	62.8	61.9	65.2	64.5	59.9	58.8
	R11 (7y)	307	55.4	52.2	54.6	54.2	60.6	50.1	59.7	62.8	65.5	68.1	63.5	58.8
	R12 (5y)	307	56.6	50.8	55.6	54.4	60.3	50.3	60.2	63.2	62.1	62.4	63.5	58.1
	R10 (7y)	307	54.0	50.5	55.0	52.3	62.0	56.6	56.9	61.1	64.6	59.9	62.6	57.8
	R17 (5y)	307	55.0	51.7	56.1	53.1	61.3	50.4	61.5	59.0	66.1	63.1	56.9	57.7
	R27 (2y)	307	56.9	53.1	57.4	55.6	59.3	52.2	54.9	60.0	63.4	59.6	58.8	57.4
	R20 (3y)	307	55.8	51.8	55.3	53.4	58.5	51.2	59.6	61.2	63.8	61.3	57.5	57.2
	R01 (6y)	307	53.5	51.8	54.0	52.9	57.1	53.1	57.3	63.3	63.1	62.8	59.8	57.2
	R03 (9y)	307	54.7	49.5	56.2	53.1	58.5	51.1	54.2	59.8	63.4	59.6	61.1	56.5
	R08 (7y)	307	50.8	47.6	52.8	51.1	60.6	54.0	57.3	57.7	63.4	66.4	59.0	56.4
	R04 (8y)	307	54.9	51.6	54.0	51.1	60.7	51.6	56.6	59.9	63.4	60.2	56.0	56.4
	R06 (8y)	307	53.6	50.0	54.9	51.1	62.0	49.9	58.4	57.9	61.8	61.8	58.1	56.3
	R13 (5y)	307	53.0	49.7	54.5	52.5	57.4	54.5	52.8	59.6	61.7	62.2	58.8	56.1
	R25 (3y)	307	54.4	50.4	54.8	53.0	54.1	49.0	58.0	53.5	62.7	62.6	61.5	55.8
	R14 (5y)	307	52.5	50.1	54.3	52.0	56.2	50.1	58.0	60.0	60.5	59.4	54.9	55.3
	R02 (9y)	307	48.6	44.9	49.7	47.5	60.4	52.4	56.7	57.6	62.6	61.2	58.8	54.6
	R26 (2y)	307	52.8	48.5	53.1	52.0	53.8	50.2	53.7	56.0	60.5	61.6	57.9	54.6
	R24 (3y)	307	49.7	48.3	51.0	50.6	57.6	47.6	53.9	62.9	60.2	60.8	56.1	54.4
	R18 (4y)	307	52.6	49.8	53.7	50.5	52.7	51.3	55.6	55.9	60.8	57.5	57.0	54.3
	R07 (7y)	307	53.8	49.5	53.5	50.1	55.5	48.4	54.1	53.0	59.6	58.8	56.2	53.9
	R15 (5y)	307	53.7	50.3	53.2	51.9	54.6	49.0	54.5	54.0	60.6	56.9	53.7	53.9
	R16 (5y)	307	53.0	49.2	53.5	51.1	53.8	47.4	53.0	54.4	60.8	60.4	53.8	53.7
	R21 (3y)	307	38.1	35.8	39.2	36.9	55.1	45.0	56.3	53.9	56.7	59.3	53.6	48.2
	R19 (3y)	307	39.0	35.9	38.4	37.5	51.2	45.1	50.7	54.4	55.6	55.4	51.4	46.8
	Mean	307	53.6	50.2	54.2	52.1	58.2	50.8	57.1	59.2	62.5	61.4	58.5	56.2
AI+Resident	R22 (3y)	307	74.4	76.1	81.0	75.5	75.2	69.5	85.3	86.9	89.5	88.8	65.8	78.9
	R05 (8y)	307	66.9	73.9	84.6	74.1	73.2	67.3	80.0	80.9	86.4	87.3	77.9	77.5
	R10 (7y)	307	64.1	60.4	68.5	57.9	81.1	75.1	82.4	87.3	86.9	89.7	89.1	76.6
	R25 (3y)	307	69.4	62.7	78.6	62.7	77.8	70.2	76.9	79.6	82.9	88.6	81.4	75.5
	R14 (5y)	307	77.1	65.3	59.8	69.0	75.8	69.2	77.1	78.8	88.6	86.1	79.5	75.1
	R26 (2y)	307	74.5	63.8	74.4	62.7	71.2	66.6	79.0	79.3	82.0	85.4	80.2	74.5
	R13 (5y)	307	79.7	59.3	64.3	68.7	71.8	68.8	76.8	79.6	82.5	89.5	77.5	74.4
	R07 (7y)	307	78.1	56.7	67.9	64.4	77.5	67.0	76.7	75.8	81.8	88.6	81.9	74.2
	R18 (4y)	307	61.8	73.6	59.6	63.4	73.1	64.3	82.5	77.1	89.8	87.6	79.3	73.8
	R01 (6y)	307	74.3	61.4	59.6	64.5	71.7	68.5	77.0	80.0	85.2	87.9	79.3	73.6
	R11 (7y)	307	67.0	77.6	69.0	64.6	76.3	63.3	79.7	78.1	83.6	67.2	79.1	73.2

(continued)

Group	Reader	n	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	T11	Overall
	R12 (5y)	307	62.9	74.0	63.2	68.7	75.9	64.7	75.6	78.0	82.8	85.0	74.9	73.2
	R08 (7y)	307	65.0	57.9	59.8	71.7	75.8	68.9	82.0	76.8	84.6	88.7	74.1	73.2
	R06 (8y)	307	67.1	58.7	61.5	60.9	77.0	68.1	81.8	81.4	83.5	86.9	77.3	73.1
	R02 (9y)	307	61.6	54.8	73.5	58.6	76.3	73.4	78.0	77.7	82.6	86.5	81.4	73.1
	R23 (3y)	307	64.1	62.0	64.7	59.9	72.3	67.6	77.9	79.0	84.3	88.2	78.1	72.6
	R04 (8y)	307	74.3	75.7	72.6	64.1	67.1	61.9	70.9	74.9	77.5	78.6	70.2	71.6
	R24 (3y)	307	60.9	54.4	60.3	53.7	75.2	66.9	84.3	81.1	85.2	87.9	76.6	71.5
	R27 (2y)	307	71.3	60.4	73.0	70.0	68.7	64.4	76.6	78.4	62.4	82.9	77.0	71.4
	R20 (3y)	307	67.6	61.0	65.6	57.3	73.6	65.0	79.9	76.3	81.5	82.4	75.1	71.4
	R03 (9y)	307	64.8	66.4	72.2	67.9	69.0	62.3	72.2	74.3	79.4	80.2	77.0	71.4
	R16 (5y)	307	58.1	61.0	58.3	59.7	74.4	65.0	77.3	78.2	82.3	86.6	81.4	71.1
	R09 (7y)	307	76.5	60.8	67.1	73.1	65.5	63.9	61.3	73.0	76.5	82.0	73.8	70.3
	R15 (5y)	307	69.2	58.4	73.1	65.0	67.1	60.2	71.1	70.9	78.4	79.4	70.6	69.4
	R21 (3y)	307	41.7	55.1	57.2	45.5	76.4	66.0	81.2	76.3	96.4	85.9	79.8	69.2
	R17 (5y)	307	66.8	58.2	69.8	60.7	72.9	59.8	70.9	70.8	74.9	78.5	70.3	68.5
	R19 (3y)	307	48.1	48.8	49.9	44.6	70.0	62.1	75.4	75.6	77.8	81.5	75.6	64.5
	Mean	307	66.9	62.9	67.0	63.3	73.4	66.3	77.4	78.0	82.6	84.7	77.2	72.7

Table J11: Adjudication criteria for prospective crossover reader study

Task	Clinical task	Reference standard definition	Scoring method	Adjudication procedure
P1	Admission diagnosis	Admission diagnosis documented in the medical record at the time of admission	Correct / Incorrect	Direct comparison with the admission diagnosis documented in the admission note; semantically equivalent diagnoses accepted after review by two senior surgeons
P2	Preoperative diagnosis	Preoperative diagnosis documented in the preoperative assessment immediately before surgery	Correct / Incorrect	Direct comparison with the preoperative diagnosis documented in the preoperative assessment; semantically equivalent diagnoses accepted after review by two senior surgeons
P3	Surgical planning	No single ground truth; quality evaluated by expert committee across five prespecified dimensions: approach selection, fixation strategy, risk preparedness, patient-specific adaptation, and plan feasibility	Weighted composite score (1–10 scale per dimension)	Each case was independently scored by at least three of fifteen senior orthopaedic surgeons (≥ 15 years of experience) who were blinded to study condition, and the final score was calculated as the mean across raters
P4	Intraoperative diagnosis	Final intraoperative diagnosis documented in the operative report	Correct / Incorrect	Direct comparison with the operative diagnosis in the surgical record; borderline or discordant cases reviewed by two senior surgeons

Table J12: Evaluation criteria for surgical planning quality assessment			
Evaluation dimension	Weight	Score range	Scoring criteria
Surgical approach selection appropriateness	20%	1–10	Whether the selected surgical approach is appropriate for the injury type and anatomical characteristics, and whether adequate exposure is balanced with soft-tissue protection
Fixation or implant strategy suitability	25%	1–10	Whether fixation method and implant selection match fracture type or joint pathology, and whether biomechanical principles are appropriately considered
Intraoperative risk anticipation and preparedness	20%	1–10	Whether potential intraoperative risks are adequately identified, such as neurovascular injury, osteoporosis, or implant loosening, and whether reasonable contingency plans are proposed
Patient-specific adaptation	20%	1–10	Whether the plan sufficiently considers patient-specific factors including age, bone quality, comorbidities, and functional requirements
Plan completeness and feasibility	15%	1–10	Whether the surgical plan is comprehensive, including positioning, procedural steps, and instrument preparation, and whether it is practically executable

Table J13: Score interpretation criteria for surgical planning evaluation

Score range	Grade	Description
9–10	Excellent	The plan reaches the highest level of expert consensus in this dimension, with no substantial room for improvement
7–8	Good	The plan is generally appropriate, with minor areas for optimization that do not affect surgical safety or expected outcomes
5–6	Fair	The plan is feasible but contains notable limitations that may affect surgical outcomes or increase procedural risk
3–4	Poor	The plan contains significant deficiencies and requires substantial revision before clinical implementation
1–2	Unacceptable	The plan contains fundamental errors that may lead to serious adverse outcomes

Table J14: NASA-TLX physician workload questionnaire

Category	Dimension	Question	Response scale
Workload assessment	Mental demand	Level of cognitive effort required for clinical decision making, information processing and task execution	0–100 visual analogue scale, higher scores indicate greater demand
Workload assessment	Physical demand	Level of physical effort required during clinical work	0–100 visual analogue scale, higher scores indicate greater demand
Workload assessment	Temporal demand	Perceived time pressure during clinical workflow	0–100 visual analogue scale, higher scores indicate greater demand
Workload assessment	Performance	Self-rated performance in completing clinical tasks	0–100 visual analogue scale, higher scores indicate better performance
Workload assessment	Effort	Overall effort required to achieve performance level	0–100 visual analogue scale, higher scores indicate greater effort
Workload assessment	Frustration	Level of stress, irritation or frustration during clinical work	0–100 visual analogue scale, higher scores indicate greater frustration

Table J15: Physician documentation workload questionnaire

Category	Clinical task	Question	Response options
Documentation load	work- Admission note	Time spent completing admission documentation	Not applicable; 0–3 min; 3–6 min; 6–9 min; 9–12 min; 12–15 min; >15 min
Documentation load	work- Progress note	Time spent completing daily progress documentation	Not applicable; 0–3 min; 3–6 min; 6–9 min; 9–12 min; 12–15 min; >15 min
Documentation load	work- Order entry	Time spent entering medical orders	Not applicable; 0–3 min; 3–6 min; 6–9 min; 9–12 min; 12–15 min; >15 min
Documentation load	work- Discharge summary	Time spent completing discharge documentation	Not applicable; 0–3 min; 3–6 min; 6–9 min; 9–12 min; 12–15 min; >15 min

Table J16: Nursing workflow efficiency questionnaire

Category	Clinical task	Question	Response options
Documentation efficiency	Nursing documentation	Time spent completing nursing documentation	0–3 min; 3–6 min; 6–9 min; 9–12 min; 12–15 min; >15 min
Documentation burden	Perceived workload	Perceived burden of nursing documentation	Very low; Low; Moderate; High; Very high
Communication efficiency	Patient response	Average time required to respond to patient or family questions	0–2 min; 2–4 min; 4–6 min; >6 min

Table J17: Patient overall satisfaction questionnaire

Category	Dimension	Question	Response scale
Clinical quality	Medical expertise	Satisfaction with diagnostic and treatment quality	1–5 Likert scale, very satisfied to very satisfied
Care experience	Staff attitude	Satisfaction with physician and nurse communication and care	1–5 Likert scale, very dissatisfied to very satisfied
Workflow efficiency	Care process	Satisfaction with admission, examination, surgery and discharge workflow	1–5 Likert scale, very dissatisfied to very satisfied
Information delivery	Information clarity	Satisfaction with explanation of disease and treatment plan	1–5 Likert scale, very dissatisfied to very satisfied

Table J18: Doctor–patient communication quality questionnaire

Category	Dimension	Question	Response scale
Communication quality	Listening	Physician attentiveness to patient concerns	1–5 Likert scale, very dissatisfied to very satisfied
Communication quality	Explanation clarity	Clarity of explanation regarding diagnosis and treatment	1–5 Likert scale, very dissatisfied to very satisfied
Communication quality	Responsiveness	Physician responsiveness to patient questions	1–5 Likert scale, very dissatisfied to very satisfied
Communication quality	Respect	Respect for patient preferences and decisions	1–5 Likert scale, very dissatisfied to very satisfied
Communication quality	Information completeness	Adequacy of information provided including risks and alternatives	1–5 Likert scale, very dissatisfied to very satisfied

Table J19: Health information accessibility questionnaire

Category	Dimension	Question	Response scale
Information access	Test results	Ease of obtaining laboratory and imaging results	1–5 Likert scale, very difficult to very easy
Information understanding	under-Medical terminology	Ease of understanding medical terminology	1–5 Likert scale, very difficult to very easy
Information availability	Educational materials	Availability of disease-related educational materials	1–5 Likert scale, very difficult to very easy
Communication accessibility	Staff consultation	Ease of consulting medical staff	1–5 Likert scale, very difficult to very easy
Post-discharge guidance	Rehabilitation instructions	Clarity of post-discharge rehabilitation guidance	1–5 Likert scale, very difficult to very easy

1798

Appendix K Case Study

1799

K.1 Full-Pathway Cases

1800

1801

1802

1803

1804

1805

This subsection presents six representative full-pathway cases illustrating key reasoning behaviours of OrthoPilot along the clinical trajectory. In each figure, every dot in the evidence trajectory denotes one real tool call: purple dots indicate in-hospital tools and blue dots indicate external tools. Pale-yellow trajectory cards highlight key turning points in the reasoning process, such as evidence anchoring, fallback retrieval, guideline confirmation or risk-control decisions. The first four cases (Figs. K9–K12) demonstrate successful reasoning patterns across rehabilitation planning and surgical planning; the final two cases (Figs. K13–K14) illustrate characteristic failure modes in diagnostic output generation.

1806

1807

1808

1809

- **Rotator cuff repair rehabilitation** (Fig. K9): OrthoPilot combines in-hospital structural evidence with external guideline and literature retrieval to generate a postoperative rehabilitation plan. The evidence trajectory shows real tool calls with pale-yellow cards highlighting key turning points, including evidence anchoring, fallback retrieval, guideline confirmation and risk-control decisions.

1810

1811

1812

1813

- **L3 vertebral compression fracture after vertebroplasty** (Fig. K10): OrthoPilot identifies unreliable knowledge-graph outputs, redirects evidence gathering to patient-case retrieval and guideline search, and generates a rehabilitation plan for postoperative lower-limb motor dysfunction. This case illustrates the system’s ability to detect and recover from an unreliable evidence source mid-trajectory.

1814

1815

1816

- **Total hip arthroplasty (THA) rehabilitation** (Fig. K11): Score-informed patient-case retrieval drives staged rehabilitation planning after total hip arthroplasty. The evidence trajectory summarises real tool calls and reasoning milestones, with pale-yellow cards marking key evidence transitions and clinically decisive steps.

1817

1818

1819

1820

- **Scapular osteochondroma surgical planning** (Fig. K12): A rare anatomical tumour case demonstrating multi-tool reasoning, including recognition of unreliable structured knowledge, retrieval of analogous in-hospital cases and targeted external literature search. Pale-yellow trajectory cards mark key turning points across the multi-source evidence chain.

1821

1822

1823


- **Ankle admission-diagnosis validation — failure case** (Fig. K13): A reasoning-output mismatch in which the evidence trajectory supports the correct diagnostic judgement, but the final answer fails to provide the required validity decision. Pale-yellow cards highlight key turning points and the terminal failure point.


1824

1825

1826

- **Knee arthropathy diagnostic contradiction — failure case** (Fig. K14): A think-answer contradiction in which the reasoning trajectory rejects one diagnosis but the final output reports it. Pale-yellow cards mark the decisive evidence steps leading to the contradiction and the terminal failure point.









Fig. K11: Total hip arthroplasty rehabilitation.






Fig. K13: Ankle admission-diagnosis validation (failure case).




Fig. K14: Knee arthropathy diagnostic contradiction (failure case).

1827 K.2 OrthoBench Cases

1828 We present eleven representative case comparisons spanning all eleven OrthoBench tasks (T1–T11), covering both
1829 closed-ended diagnostic tasks (T1–T4) and open-ended management tasks (T5–T11). For each task, two representative
1830 patients are shown side by side. CHEESE is compared against DeepSeek-R1, GPT-5.1 and MedGemma-27B. High-
1831 lighted annotations in each figure identify key correct elements captured by CHEESE and the major error patterns in
1832 baseline responses. Together, the eleven figures provide a qualitative account of how model outputs differ across the
1833 full care pathway, complementing the quantitative results in the main text.

- 1834 • Fig. K15: Admission diagnosis (T1).
- 1835 • Fig. K16: Preoperative diagnosis (T2).
- 1836 • Fig. K17: Intraoperative diagnosis (T3).
- 1837 • Fig. K18: Discharge diagnosis (T4).
- 1838 • Fig. K19: Perioperative assessment (T5).
- 1839 • Fig. K20: Surgical planning (T6).
- 1840 • Fig. K21: Preoperative orders (T7).
- 1841 • Fig. K22: Postoperative orders (T8).
- 1842 • Fig. K23: Discharge summary (T9).
- 1843 • Fig. K24: Rehabilitation planning (T10).
- 1844 • Fig. K25: Multidisciplinary consultation (T11).

1845 K.3 Application Cases

1846 This subsection presents two full-deployment application cases demonstrating OrthoPilot’s end-to-end clinical workflow
1847 in a real hospital setting. Unlike the benchmark cases, these examples trace the full agent session from initial patient
1848 context loading through multi-turn tool use, evidence synthesis and final structured output, reflecting the system’s
1849 actual deployment interface.

- 1850 • **Case A: High-risk hip fracture** (Figs. K26–K29): A four-page walkthrough of a complex perioperative hip
1851 fracture episode, showing the complete agent session from evidence retrieval to final clinical outputs.
 - 1852 – Page 1 (Fig. K26): Patient-specific EHR evidence retrieval and the coordination between the main agent and
1853 subagents.
 - 1854 – Page 2 (Fig. K27): Tool marketplace configuration and subagent dispatch, with intermediate reasoning summaries
1855 produced after each tool return.
 - 1856 – Page 3 (Fig. K28): Patient-specific evidence synthesis and structured recommendation generation across
1857 perioperative care.
 - 1858 – Page 4 (Fig. K29): Final outputs organised for clinical review, patient communication and downstream
1859 documentation.
- 1860 • **Case B: Ankle arthroscopy** (Figs. K30–K31): A two-page walkthrough of an ankle arthroscopy episode,
1861 illustrating how OrthoPilot maintains a persistent session across the perioperative timeline.
 - 1862 – Page 1 (Fig. K30): Patient-specific session anchoring, EHR retrieval, imaging review and coordinated main
1863 agent/subagent reasoning.
 - 1864 – Page 2 (Fig. K31): Incorporation of intraoperative findings into the persistent session and synthesis into structured
1865 postoperative conclusions.

Admission diagnosis

Identify the correct admission diagnosis from the provided clinical evidence.

Prompt

Reference

CHEESE key points

Baseline errors

Case A

What is the likely admission diagnosis?

CLINICAL CONTEXT

Right proximal fibular destructive lesion after biopsy and chemotherapy. The prompt asks for the admission disease name.

REFERENCE

Malignant tumor of the fibula

CHEESE

Right proximal fibular malignant tumor.

MATCHED KEY POINTS

- matched malignant lesion
- localized to the fibula
- kept the admission-level label

DeepSeek-R1

Osteosarcoma.

MAIN ERROR POINTS

- over-specified a histologic subtype
- missed the required admission label

GPT-5.1

Osteosarcoma.

MAIN ERROR POINTS

- over-specified a histologic subtype
- did not match the reference label

MedGemma

Osteosarcoma.

MAIN ERROR POINTS

- over-specified a histologic subtype
- missed the fibular tumor label

Case B

What is the current clinical admission label?

CLINICAL CONTEXT

Post-traumatic right lower-leg fracture with stable external fixation and healed wounds. The admission was for device removal.

REFERENCE

Removal of an external fixation device

CHEESE

Healed right lower-leg fracture requiring external fixator removal.

MATCHED KEY POINTS

- captured external-fixator removal
- recognized healed postoperative state
- did not overcall infection

DeepSeek-R1

Chronic osteomyelitis.

MAIN ERROR POINTS

- mistook prior redness for infection
- missed device-removal target

GPT-5.1

Delayed union or nonunion with possible chronic osteomyelitis.

MAIN ERROR POINTS

- overcalled delayed union
- missed the planned removal admission

MedGemma

Osteomyelitis.

MAIN ERROR POINTS

- focused on infection
- ignored the external-fixator removal context

Fig. K15: Representative case comparison for admission diagnosis (T1).

Preoperative diagnosis		Prompt	Reference	CHEESE key points	Baseline errors
Infer the most appropriate diagnosis before surgery.					
Case A		What is the preoperative diagnosis?			
CLINICAL CONTEXT Large long-standing left gluteal cystic soft-tissue lesion with clear borders. The clinical preoperative label remained broad.		REFERENCE Subcutaneous mass			
CHEESE Left gluteal subcutaneous mass. MATCHED KEY POINTS • matched subcutaneous mass • preserved anatomic site • avoided over-specific imaging diagnosis		DeepSeek-R1 Ischial tuberosity bursitis. MAIN ERROR POINTS • over-interpreted imaging suggestion • missed the broad surgical diagnosis			
GPT-5.1 Left gluteal ischial bursitis cyst. MAIN ERROR POINTS • over-specific cystic diagnosis • did not match the reference label		MedGemma Ischial tuberosity bursitis. MAIN ERROR POINTS • over-specified a benign bursal entity • missed the target wording			
Case B		What is the most likely preoperative diagnosis?			
CLINICAL CONTEXT Left distal femoral intramedullary tumor with other suspicious bone lesions and immunoglobulin abnormalities.		REFERENCE Left distal femoral tumor			
CHEESE Left distal femoral malignant tumor-like lesion. MATCHED KEY POINTS • localized to distal femur • captured tumor diagnosis • matched the operative focus		DeepSeek-R1 Multiple myeloma. MAIN ERROR POINTS • shifted to systemic disease • missed the operative preoperative label			
GPT-5.1 Multiple myeloma, IgG-kappa type. MAIN ERROR POINTS • over-specific systemic diagnosis • ignored the expected lesion label		MedGemma Multiple myeloma. MAIN ERROR POINTS • focused on immunoglobulin findings • missed the distal femoral tumor label			

Fig. K16: Representative case comparison for preoperative diagnosis (T2).

Intraoperative diagnosis

Determine the diagnosis most directly supported by the operative findings.

Prompt

Reference

CHEESE key points

Baseline errors

Case A

What pathological change was found intraoperatively?

CLINICAL CONTEXT

L4-5 canal lesion resected during decompression surgery. Operative note described a yellow-ligament mass adherent to the dura.

REFERENCE

Intraspinal canal mass

CHEESE

L4-5 intraspinal yellow-ligament mass.

MATCHED KEY POINTS

- captured canal mass
- used operative location
- linked to yellow-ligament finding

DeepSeek-R1

Neurilemmoma.

MAIN ERROR POINTS

- invented a tumor type
- missed the operative generic label

GPT-5.1

Lumbar ligamentum flavum cyst.

MAIN ERROR POINTS

- over-specific cyst diagnosis
- did not match the reference wording

MedGemma

Lumbar ligamentum flavum tumor.

MAIN ERROR POINTS

- over-specified pathology
- missed the broad canal-mass label

Case B

What is the most likely intraoperative diagnosis?

CLINICAL CONTEXT

Right knee pain with osteophytes, effusion, medial compartment disease and arthroplasty-related operative findings.

REFERENCE

Degenerative disease of the knee joint

CHEESE

Right knee osteoarthritis, a degenerative knee-joint disease.

MATCHED KEY POINTS

- matched degenerative disease
- captured right knee OA
- used operative arthroplasty context

DeepSeek-R1

Knee osteoarthritis.

MAIN ERROR POINTS

- partially related term
- missed the exact degenerative-disease label

GPT-5.1

Medial compartment knee osteoarthritis.

MAIN ERROR POINTS

- over-narrowed to compartmental OA
- partial terminology mismatch

MedGemma

Knee osteoarthritis.

MAIN ERROR POINTS

- related but incomplete
- failed to match the expected wording

Fig. K17: Representative case comparison for intraoperative diagnosis (T3).

Discharge diagnosis

Summarize the diagnosis that best represents the patient at this stage.

Prompt

Reference

CHEESE key points

Baseline errors

Case A

What is the most likely final diagnosis?

CLINICAL CONTEXT

Longstanding bilateral hip pain after prior right hip arthroplasty. Imaging showed irregular femoral head with subchondral sclerosis.

REFERENCE

Avascular necrosis of the femoral head

CHEESE

Bilateral avascular necrosis of the femoral head.

MATCHED KEY POINTS

- matched femoral-head necrosis
- recognized bilateral disease
- avoided OA-only diagnosis

DeepSeek-R1

Bilateral hip osteoarthritis.

MAIN ERROR POINTS

- overcalled degenerative arthritis
- missed femoral-head necrosis

GPT-5.1

Ankylosing spondylitis involving the hips.

MAIN ERROR POINTS

- introduced unsupported systemic disease
- missed AVN

MedGemma

Hip osteoarthritis.

MAIN ERROR POINTS

- generic degenerative label
- failed to identify AVN

Case B

What is the final diagnostic label?

CLINICAL CONTEXT

Right second metatarsal benign-appearing lesion treated by curettage, grafting and fixation. Pathology suggested osteoblastoma.

REFERENCE

Benign tumor of the metatarsal bone

CHEESE

Benign tumor of the right second metatarsal bone.

MATCHED KEY POINTS

- matched benign metatarsal tumor
- preserved site
- kept the clinical-label level

DeepSeek-R1

Epithelioid osteoblastoma.

MAIN ERROR POINTS

- gave pathology subtype
- missed the expected broad label

GPT-5.1

Epithelioid osteoblastoma.

MAIN ERROR POINTS

- over-specific pathology subtype
- did not match the reference category

MedGemma

Epithelioid osteoblastoma.

MAIN ERROR POINTS

- focused on subtype
- missed benign metatarsal tumor wording

Fig. K18: Representative case comparison for discharge diagnosis (T4).

Perioperative assessment

Judge whether the patient can appropriately proceed with surgery.

Prompt

Reference

CHEESE key points

Baseline errors

Case A

Can the patient undergo surgery?

CLINICAL CONTEXT

Left ankle injury with mild abnormalities, good AOFAS score, no major systemic contraindication, and planned surgery assessment.

REFERENCE

Yes, the patient can undergo surgery.

CHEESE

Yes, the patient can undergo surgery.

MATCHED KEY POINTS

- matched yes/no decision
- recognized no absolute contraindication
- did not over-penalize mild labs

DeepSeek-R1

No immediate surgery; treat urinary infection first.

MAIN ERROR POINTS

- overestimated urinalysis abnormality
- changed the binary decision

GPT-5.1

Not recommended immediately; treat urinary infection and reassess.

MAIN ERROR POINTS

- overcalled perioperative risk
- missed overall surgical clearance

MedGemma

Not an immediate candidate unless later indicated.

MAIN ERROR POINTS

- questioned indication rather than eligibility
- missed the reference yes answer

Case B

Can the patient undergo surgery?

CLINICAL CONTEXT

Multiple left-foot fractures with intact distal circulation and no clear absolute contraindication to operative treatment.

REFERENCE

Yes, the patient can undergo surgery.

CHEESE

Yes, the patient can undergo surgery.

MATCHED KEY POINTS

- matched yes decision
- recognized fracture indication
- did not require unnecessary delay

DeepSeek-R1

Suitable for surgery, but many extra optimization steps are needed.

MAIN ERROR POINTS

- partly aligned
- over-expanded with nonessential risks

GPT-5.1

Has surgical indication, but not immediately; delay until soft tissue and labs improve.

MAIN ERROR POINTS

- converted approval into delay
- overstated relative concerns

MedGemma

Not an immediate elective-surgery candidate due to neutropenia and endocrine concerns.

MAIN ERROR POINTS

- overcalled relative lab findings
- missed operative feasibility

Fig. K19: Representative case comparison for perioperative assessment (T5).

Surgical planning		Prompt	Reference	CHEESE key points	Baseline errors
Formulate the key pre-surgical plan required for the next step of care.					
Case A		What preoperative orders should be arranged?			
CLINICAL CONTEXT Right gluteal soft-tissue tumor planned for excision after laboratory workup, pelvic MRI, skin preparation and blood preparation.		REFERENCE Level-I nursing, NPO, renal/liver/coagulation/infection tests, skin prep, 4-unit blood prep, fluids, antibiotics, planned gluteal tumor excision.			
CHEESE Preoperative tests, NPO, skin prep, 4-unit blood preparation, saline infusion and gluteal tumor excision under anesthesia. MATCHED KEY POINTS • captured core tests • included NPO, skin prep and blood prep • included planned procedure		DeepSeek-R1 Adds extensive anesthesia, catheter, antibiotic and monitoring orders. MAIN ERROR POINTS • missed exact order set • added many non-reference items			
GPT-5.1 Adds staging, biopsy, MDT and enhanced tumor workup. MAIN ERROR POINTS • overcomplicated beyond orders • omitted several reference order details		MedGemma Focuses on monitoring, renal function and bowel preparation. MAIN ERROR POINTS • missed key reference items • added irrelevant bowel preparation			
Case B		What is a reasonable preoperative order plan?			
CLINICAL CONTEXT Right third toe fracture after fixation one year earlier. X-ray showed healed fracture and satisfactory hardware position; hardware removal was planned.		REFERENCE Regular diet, NPO from 8 PM, urinalysis/sediment, ECG, APTT, skin preparation, no special preoperative medication, hardware-removal surgery.			
CHEESE Regular diet, NPO before surgery, APTT/coagulation assessment, skin preparation and no special drug preparation. MATCHED KEY POINTS • captured NPO and diet • included APTT/coagulation • recognized no special drug prep		DeepSeek-R1 Adds respiratory, renal and anesthesia consultations plus clindamycin. MAIN ERROR POINTS • overcalled chest and renal findings • added unnecessary antibiotic plan			
GPT-5.1 Detailed elective minor-surgery protocol with many added tests and antibiotics. MAIN ERROR POINTS • too broad for target answer • missed concise order-set format		MedGemma Repeats broad preoperative checks, chest film and prophylactic antibiotics. MAIN ERROR POINTS • added non-reference items • less focused on the expected order list			

Fig. K20: Representative case comparison for surgical planning (T6).

Preoperative orders		Prompt	Reference	CHEESE key points	Baseline errors
Specify the essential operative or procedural steps expected for the planned intervention.					
Case A		Describe the operative steps and technical points.			
CLINICAL CONTEXT Elderly woman with recurrent lumbar symptoms after prior fusion and fixation. The prompt asks for posterior decompression, interbody fusion and fixation steps.		REFERENCE Prone position, midline exposure, pedicle screws, decompression, root canal enlargement, disc removal, endplate preparation, cage insertion, rods, fluoroscopy, irrigation, drain and closure.			
CHEESE Prone positioning, midline exposure, decompression, nerve-root canal enlargement, disc removal, cage placement, pedicle screws, rods, fluoroscopy and closure. MATCHED KEY POINTS - covered key operative sequence - included decompression and root protection - included cage and fixation		DeepSeek-R1 Detailed revision strategy with extra extension fixation and osteoporosis planning. MAIN ERROR POINTS - added non-reference revision strategy - missed concise step order			
GPT-5.1 Very broad PLIF/TLIF tutorial with many optional details. MAIN ERROR POINTS - overly generic and lengthy - less faithful to the target operation note		MedGemma Includes laminectomy, discectomy and fusion but also says to open the dural sac. MAIN ERROR POINTS - dangerous inaccurate step - added non-reference maneuvers			
Case B		Describe the steps and precautions for external-fixator removal.			
CLINICAL CONTEXT Right tibia/fibula fracture after external fixation with callus formation. The operation was removal of the lower-leg external fixation device.		REFERENCE Supine position, sterile preparation, remove screws, cut ring wires, remove the ring external frame, irrigate, close and dress.			
CHEESE Supine positioning, sterile preparation, removal of pins and frame components, wound irrigation, layered closure and sterile dressing. MATCHED KEY POINTS - captured core removal workflow - included irrigation and closure - preserved uncomplicated procedure		DeepSeek-R1 Adds extensive assessment, screw extraction, cultures, antibiotics and rehabilitation. MAIN ERROR POINTS - over-expanded beyond operative steps - missed ring-wire cutting detail			
GPT-5.1 Comprehensive Schanz-pin removal protocol with pin-tract management. MAIN ERROR POINTS - not tailored to ring-frame wording - too much non-reference management		MedGemma General external-fixator removal, mostly pin and frame removal. MAIN ERROR POINTS - generic and incomplete - missed key ring-frame sequence			

Fig. K21: Representative case comparison for preoperative orders (T7).

Postoperative orders

Select appropriate orders, medications and follow-up management after surgery.

Prompt

Reference

CHEESE key points

Baseline errors

Case A

Which postoperative orders are most likely needed on Oct 27?

CLINICAL CONTEXT

Polytrauma with calcaneal and fibular fractures and degloving wound treated by debridement, VSD drainage, antibiotics and rehabilitation.

REFERENCE

Level-II nursing, regular diet, lower-limb X-ray and glucose monitoring, low-frequency pulse therapy, joint mobilization, functional exercise, drainage care, diosmin and celecoxib.

CHEESE

Level-II nursing, regular high-protein diet, rehabilitation training, drainage management and analgesic/vascular medication.

MATCHED KEY POINTS

- matched nursing level and diet
- included rehab and drainage
- included key medication direction

DeepSeek-R1

A broad recovery plan with antibiotics and monitoring.

MAIN ERROR POINTS

- missed exact order additions
- failed to include core rehab modalities

GPT-5.1

General postoperative care and follow-up recommendations.

MAIN ERROR POINTS

- too generic
- missed X-ray, glucose and therapy items

MedGemma

General discharge and wound-care advice.

MAIN ERROR POINTS

- not formatted as orders
- omitted key order items

Case B

What postoperative order arrangement is needed?

CLINICAL CONTEXT

Left humeral pathological fracture after curettage, grafting and internal fixation, with drain, anemia/hypoalbuminemia and follow-up imaging.

REFERENCE

Level-II nursing, regular diet, left humerus X-ray, urgent renal tests, glucose monitoring, single-lumen drain care, major dressing change, IV infusion, flurbiprofen, albumin, dexamethasone, mecobalamin and escin.

CHEESE

Level-II nursing, regular high-protein diet, left humerus X-ray, urgent renal tests, glucose monitoring, albumin, flurbiprofen, dexamethasone, drain care, mecobalamin and escin.

MATCHED KEY POINTS

- captured core orders
- included drain and imaging
- included key medications

DeepSeek-R1

General postoperative management with lab monitoring and infection control.

MAIN ERROR POINTS

- missed exact medication/order set
- not aligned with order-list style

GPT-5.1

Adds transfusion thresholds, oncology follow-up and broad monitoring.

MAIN ERROR POINTS

- over-expanded beyond reference
- omitted some exact orders

MedGemma

Provides routine postoperative care and rehabilitation advice.

MAIN ERROR POINTS

- less specific to requested orders
- missed multiple named medications

Fig. K22: Representative case comparison for postoperative orders (T8).

50

Discharge summary

Summarize admission, treatment course, operation and discharge status.

Prompt

Reference

CHEESE key points

Baseline errors

Case A

Generate the discharge-course summary.

CLINICAL CONTEXT

Right distal femoral osteosarcoma treated by tumor-segment resection and prosthetic replacement under general anesthesia.

REFERENCE

Admission for right distal femoral osteosarcoma, preoperative workup, no contraindication, smooth tumor-segment resection and prosthetic replacement, good implant position, stable recovery and regular follow-up.

CHEESE

Admission for right distal femoral osteosarcoma; workup excluded contraindications; tumor resection and prosthesis replacement were uneventful; imaging showed good implant position; recovery met discharge criteria.

MATCHED KEY POINTS

- captured disease and operation
- included no complications
- included implant position and discharge status

DeepSeek-R1

Chronological summary with many lab and imaging details.

MAIN ERROR POINTS

- too verbose
- less focused on discharge-course essentials

GPT-5.1

Detailed structured diagnosis, workup, operation and labs.

MAIN ERROR POINTS

- partially correct
- overloaded with nonessential details

MedGemma

Broad narrative summary with extensive lab repetition.

MAIN ERROR POINTS

- less concise
- underemphasized discharge readiness

Case B

Generate the discharge-course summary.

CLINICAL CONTEXT

Lumbar scoliosis with canal stenosis treated by anterior lumbar discectomy, decompression and interbody fusion across multiple levels.

REFERENCE

Admission for lumbar scoliosis with canal stenosis, preoperative ECG/CT/X-ray workup, diagnosis of scoliosis with disc protrusion and stenosis, smooth anterior decompression and fusion, normal wound healing and discharge.

CHEESE

Lumbar scoliosis with degenerative stenosis was evaluated by ECG, CT and X-ray; anterior discectomy, decompression and interbody fusion were performed smoothly; postoperative imaging showed fixation in place and improved scoliosis; recovery allowed discharge.

MATCHED KEY POINTS

- captured diagnosis and workup
- included anterior fusion operation
- included recovery and discharge

DeepSeek-R1

Long multi-section summary of diagnoses, operation and imaging.

MAIN ERROR POINTS

- too detailed for discharge abstract
- less concise than reference

GPT-5.1

Structured but lengthy summary with numerous imaging and lab details.

MAIN ERROR POINTS

- partially aligned
- overemphasized ancillary details

MedGemma

Broad course summary with postoperative labs and comorbid findings.

MAIN ERROR POINTS

- less focused on discharge outcome
- added nonessential data

Fig. K23: Representative case comparison for discharge summary (T9).

Rehabilitation planning

Provide a focused rehabilitation evaluation and treatment plan.

Prompt

Reference

CHEESE key points

Baseline errors

Case A

What rehabilitation assessment and plan are appropriate?

CLINICAL CONTEXT

Complex perineal malignancy after Miles operation, tumor resection, colostomy, flap transfer and debridement, with pain and hip-motion limitation.

REFERENCE

Postoperative malignant perianal/thigh tumor state with pain, movement dysfunction, hip ROM restriction, gait impairment and ADL decline; recommend systematic assessment, lower-limb joint mobilization, exercise therapy, OT/ADL training, energy conservation, fall prevention and follow-up.

CHEESE

Postoperative status with wound pain, hip ROM restriction, motor dysfunction and ADL decline; recommends systematic rehabilitation assessment, joint mobilization, exercise therapy, OT/ADL training, fall prevention and follow-up.

MATCHED KEY POINTS

- captured main rehab impairments
- included ADL/OT component
- included fall prevention and follow-up

DeepSeek-R1

General rehabilitation plan with wound care and infection monitoring.

MAIN ERROR POINTS

- missed several functional targets
- less aligned with prescribed modalities

GPT-5.1

Broad oncology rehabilitation and wound-care plan.

MAIN ERROR POINTS

- partially related
- less specific to hip ROM and ADL goals

MedGemma

General postop rehab advice with strengthening and mobility.

MAIN ERROR POINTS

- generic
- omitted key structured assessment details

Case B

What rehabilitation assessment and plan are appropriate?

CLINICAL CONTEXT

Elderly patient one day after right knee revision for prosthetic joint infection, with pain and PROM limited to 0-30 degrees.

REFERENCE

Right knee revision postoperative state, NRS 2-3/10, limited PROM 0-30 degrees; comprehensive assessment, TENS around incision, joint mobilization, lower-limb exercise, balance and gait training, twice daily for five days per week, fall prevention and follow-up.

CHEESE

Right knee revision postoperative dysfunction with pain and ROM limitation; recommends assessment, joint mobilization, lower-limb exercise therapy, ADL-oriented therapy, pain monitoring and fall prevention.

MATCHED KEY POINTS

- captured pain and ROM limitation
- included joint mobilization and exercise
- included treatment frequency and safety monitoring

DeepSeek-R1

General revision-knee rehabilitation with infection and implant precautions.

MAIN ERROR POINTS

- missed TENS and exact frequency
- less focused on reference plan

GPT-5.1

Detailed phased rehabilitation plan.

MAIN ERROR POINTS

- overly broad
- omitted some required modalities

MedGemma

General strengthening and range-of-motion rehabilitation.

MAIN ERROR POINTS

- generic rehab advice
- missed TENS and schedule details

Fig. K24: Representative case comparison for rehabilitation planning (T10).

Baseline errors

Provide a coordinated rehabilitation-oriented recommendation for early recovery.

What early postoperative rehabilitation prescription is needed?

REFERENCE

One day after left shoulder arthroscopic acromioplasty for chronic shoulder pain and restricted motion, with mild postoperative pain.

Left shoulder arthroscopic acromioplasty postoperative day 1; NRS 2-3; prescribe two rehab assessments, two low-frequency pulse electrical stimulation sessions around the wound, joint mobilization and exercise therapy for the left shoulder, twice daily with gentle technique and follow-up.

DeepSeek-R1

Postoperative left shoulder arthroscopy state with limited motion and NRS 2-3; recommends pain control, joint mobilization, exercise therapy and follow-up.

Broad shoulder rehabilitation program with ROM and strengthening.

MAIN ERROR POINTS

- captured postoperative shoulder limitation
- included joint mobilization and exercise
- recognized mild pain

- missed low-frequency stimulation
- did not specify twice-daily prescription

MedGemma

Detailed phased shoulder rehabilitation protocol.

Generic frozen-shoulder postoperative rehabilitation.

MAIN ERROR POINTS

- over-detailed and generic
- missed exact order frequency

- missed acromioplasty-specific order
- omitted stimulation sessions

What early postoperative rehabilitation prescription is needed?

REFERENCE

One day after right total hip arthroplasty with osteotomy for steroid-related femoral-head necrosis, with PROM 0-30 degrees and ADL limitation.

Right total hip arthroplasty with osteotomy postoperative day 1, NRS 2-3, PROM 0-30 degrees, pain, motor dysfunction, gait/ balance and ADL limitation; prescribe rehab assessment, TENS to right hip, joint mobilization, right-limb exercise therapy, OT/ADL training twice daily for five days per week, fatigue control, fall prevention and follow-up.

DeepSeek-R1

Right total hip arthroplasty postoperative pain, ROM restriction and ADL decline; recommends rehabilitation assessment, joint mobilization, exercise therapy, ADL-oriented OT, pain and fatigue monitoring.

General hip replacement rehabilitation plan.

MAIN ERROR POINTS

- captured pain, ROM and ADL limits
- included mobilization, exercise and OT
- included safety monitoring

- missed TENS and exact schedule
- less specific to reference diagnosis

MedGemma

Detailed staged THA rehabilitation protocol.


General postoperative hip-strengthening advice.

MAIN ERROR POINTS

- too generic/staged
- did not match exact early prescription

- omitted ADL and TENS details
- missed treatment frequency

Fig. K25: Representative case comparison for multidisciplinary consultation (T11).







Fig. K27: High-risk hip fracture care, page 2 of 4.








Fig. K29: High-risk hip fracture care, page 4 of 4.





1866 Appendix L Prompts

1867 L.1 Agent Prompts

1868 The agent prompts define the control logic of OrthoPilot. They specify the system role of the main agent, the summarization behavior, the worker behavior of delegated sub-agents, and the tool-facing interaction pattern. These prompts standardize stepwise reasoning, controlled tool use, and concise answer synthesis across the orthopaedic care pathway. 1870 The corresponding prompt figures are shown in Figs. L32–L36.

1872 L.2 QA generation prompt

1873 The supervised fine-tuning data for CHEESE is constructed at three complementary granularities: single-turn QA-pair generation for closed-ended diagnostic tasks, single-turn cold-start generation for the remaining open-ended tasks, and multi-turn long-horizon agentic trajectory generation along the full inpatient care pathway. Across all three regimes, the prompts are strictly *student-facing*: the gold-standard answer is never shown inside the prompt that the model sees, so each training example forces the model to infer the target from the patient evidence alone rather than from any answer-revealing cue.

1879 *Single-turn QA-pair prompts.*

1880 QA-pair generation is reserved for the two diagnostic tasks for which a single standardised disease name is the natural training target – admission diagnosis ($\mathcal{N}_{\text{AdmDx}}$, Task 1) and preoperative diagnosis ($\mathcal{N}_{\text{PreopDx}}$, Task 2). For each of these two tasks, the prompt provides (i) a system role establishing the orthopaedic-clinician persona, (ii) a user message that supplies only the patient evidence visible at that pathway stage, and (iii) an output specification that constrains the response to a single standardised diagnosis with no reasoning, candidate list or explanation. The gold diagnosis itself is harvested directly from the structured EHR field (admission_diagnosis or the surgical record’s preoperative diagnosis) during data curation and used as the SFT target; it never appears inside the prompt. The two QA-pair prompts are shown in Figs. L37–L38. All remaining pathway tasks ($\mathcal{N}_{\text{IntraopDx}}$, $\mathcal{N}_{\text{DischDx}}$ and the management tasks $\mathcal{N}_{\text{PeriAssess}}$, $\mathcal{N}_{\text{SurgPlan}}$, $\mathcal{N}_{\text{PreopOrder}}$, $\mathcal{N}_{\text{PostopOrder}}$, $\mathcal{N}_{\text{DischSum}}$, $\mathcal{N}_{\text{Rehab}}$, $\mathcal{N}_{\text{MDT}}$) do not require teacher-generated QA pairs because their gold answers can be *extracted directly* from the structured fields of the patient record – the Intraoperative diagnosis, discharge diagnosis, operative-record narrative, orders, discharge summary, rehabilitation prescription, and consultation reports already exist in the EHR and are used verbatim as supervision targets.

1892 *Single-turn cold-start prompts.*

1893 For the tasks whose gold answer is long-form and benefits from explicit clinical reasoning, we additionally synthesise *cold-start* training records in which the model is required to emit a $\langle \text{think} \rangle$ -wrapped reasoning trace before the final answer. The cold-start prompt set covers five representative tasks: Intraoperative diagnosis ($\mathcal{N}_{\text{IntraopDx}}$, Task 3) and discharge diagnosis ($\mathcal{N}_{\text{DischDx}}$, Task 4) on the diagnostic side, and perioperative operability assessment ($\mathcal{N}_{\text{PeriAssess}}$, Task 5), surgical planning ($\mathcal{N}_{\text{SurgPlan}}$, Task 6), and discharge summary ($\mathcal{N}_{\text{DischSum}}$, Task 9) on the open-ended management side. Each prompt is student-facing, with the gold answer never inserted into the input, and specifies (i) a system role; (ii) the patient evidence available at the corresponding pathway stage; (iii) an output specification requiring at least five numbered reasoning steps inside a single $\langle \text{think} \rangle \dots \langle / \text{think} \rangle$ block, followed by a structured final answer on a new line; and (iv) task-specific structural constraints on the final answer (e.g. four-dimensional coverage for operability, a seven-section operative-step structure for surgical planning, a seven-section time-anonymised structure for the discharge summary). The five cold-start prompts are shown in Figs. L39–L43.

1904 *Long-horizon agentic trajectory prompts.*

1905 The long-horizon (LH) prompts orchestrate multi-turn, tool-augmented trajectory synthesis along the full inpatient care pathway. Unlike the QA-pair and cold-start prompts – which treat each task in isolation – the LH prompts preserve the evidence trajectory accumulated by all preceding pathway nodes and require the agent to reason within the temporal-causality envelope of real inpatient care; the gold answer is still excluded from the input, but the agent may now invoke tools from the dual evidence worlds to acquire evidence and to verify its hypotheses. Each of the 11 LH prompts shares a common system role establishing the residency-assistant persona, the tool inventory of both the in-hospital information world and the external knowledge world, and the canonical pathway sequence $\mathcal{N}_{\text{AdmDx}} \rightarrow (\mathcal{N}_{\text{PreopOrder}} \leftrightarrow \mathcal{N}_{\text{PeriAssess}})^* \rightarrow \mathcal{N}_{\text{PreopDx}} \rightarrow \mathcal{N}_{\text{MDT}} \rightarrow \mathcal{N}_{\text{SurgPlan}} \rightarrow \mathcal{N}_{\text{IntraopDx}} \rightarrow \mathcal{N}_{\text{PostopOrder}} \rightarrow \mathcal{N}_{\text{DischDx}} \rightarrow \mathcal{N}_{\text{DischSum}} \rightarrow \mathcal{N}_{\text{Rehab}}$. Each prompt additionally specifies (i) the pathway position and the in-hospital tool subset that is unlocked at that position, (ii) either a mandatory tool-call protocol (for tasks $\mathcal{N}_{\text{AdmDx}}$, $\mathcal{N}_{\text{IntraopDx}}$, $\mathcal{N}_{\text{PeriAssess}}$, $\mathcal{N}_{\text{PostopOrder}}$) that fixes the round-by-round tool-call sequence and total round budget, or recommended tool-use guidance for the remaining nodes, (iii) an output specification requiring a $\langle \text{think} \rangle$ -wrapped reasoning trace followed by a structured final answer, and (iv) quality guardrails enforcing evidence-grounding, cross-source verification, and strict temporal ordering, together with the trajectory-level rejection-sampling acceptance criterion under which only fully-correct paths are retained. The complete LH prompt set is shown in Figs. L44–L54.

1920 **L.3 Training prompt**

1921 The training prompt defines the instruction format for full-pathway agentic reinforcement learning. It frames the model
1922 as an orthopedics resident assistant and guides evidence-based diagnostic reasoning with structured tool use. The
1923 prompt specifies diagnostic principles, tool-use policies, and XML-style output formats for either evidence gathering
1924 or final diagnosis generation. The corresponding prompt figure is shown in Fig. [L55](#).

1925 **L.4 Evaluation prompts**

1926 The evaluation prompts define how reference answers are converted into structured rubrics and how model outputs
1927 are judged against them. They cover the shared grading templates and the task-specific criteria for perioperative
1928 assessment, surgical planning, preoperative orders, postoperative orders, discharge summaries, rehabilitation planning,
1929 and multidisciplinary consultation. The corresponding prompt figures are shown in Figs. [L56–L66](#).

Prompt 1: Prompt for Main Agent System

You are OrthoPilot, an evidence-guided assistant for the full orthopaedic care pathway. Your role is to combine in-hospital evidence with external medical knowledge and support clinical decision-making with transparent reasoning.

Approach each request as a sequence of small, explicit steps. Start by identifying the main objective, break it into a short ordered plan, and work through the plan one step at a time. Revise the plan only when new evidence materially changes the situation.

Use tools only when they add needed information. Stop tool use immediately and answer directly if any of the following is true: the question is already answerable from general knowledge, you have gathered enough evidence for a complete answer, a previous sub-agent report already covers the issue, or the next query would only repeat an earlier search.

When a tool call is needed, make exactly one tool call in that response. Before the call, briefly state what is already known, what remains uncertain, and why the chosen tool is the most useful next step. Every tool request must be self-contained, with enough context to stand on its own, and must avoid placeholder values, speculative assumptions, or vague search phrasing.

If the user provides image attachments, extract the full absolute file path from the attachment metadata and use the image understanding tool with that exact path. Do not shorten the path to a filename.

After issuing a tool call, stop immediately. Do not predict the result, do not continue reasoning past the call, and do not stack multiple tool invocations in a single response. Once sufficient evidence has been gathered, answer in plain text without further tool use.

Unless the user asks otherwise, reply in the same language as the user. Do not mention tool names in narrative text. Do not use nonexistent tools. Do not end with vague follow-up offers or empty questions.

Every response must begin with ``<think>`` and place all internal reasoning inside ``<think>...</think>``. After ``</think>``, output either the single tool call or the final user-facing answer.

Fig. L32: Prompt for main agent system.

Prompt 2: Prompt for Main Agent Summarization

This is the final response stage. No further tool use is allowed under any circumstances.

Produce a complete report that answers the original question using everything collected during the session. If the task ended because of a limit or failure, say so clearly. If a valid answer was already established earlier, extract that answer rather than recomputing it. If the evidence remained incomplete, provide the best supported conclusion available and mark unresolved uncertainty explicitly.

Your summary should include the full reasoning trail that matters for the final answer: the main evidence gathered, relevant intermediate findings, important partial results, contradictions, and any directly useful facts or quotations. If the task could not be fully resolved, preserve the partial findings so that a downstream reader could continue from them.

Organize the answer into clear sections with specific headings. End with a `References` section listing all useful URLs as markdown links.

All reasoning must remain inside ``<think>...</think>``. After ``</think>``, output the entire final report inside ``\boxed{...}``. Do not output a short boxed answer and a separate long answer. The boxed content must contain the full structured report.

Fig. L33: Prompt for main agent summarization.

Prompt 3: Prompt for Sub-Agent Worker System

You are the evidence-collection worker for OrthoPilot. Your job is to handle a narrowly defined subtask, gather the most relevant information, and return a precise factual report.

Break the assigned subtask into a short plan and execute it step by step. Keep the scope tight. Do not broaden the task, speculate beyond the evidence, or invent missing details.

Stop tool use as soon as any of the following occurs: one to three tool calls have already produced enough evidence, a tool returns an error or empty result, the next query would substantially duplicate an earlier one, or the answer can already be written clearly from what you have. Do not continue searching only for completeness.

When you do need a tool, make exactly one tool call in that response. Before the call, briefly state what is known, what is still missing, and why that tool is the correct next step. Every tool query must include the full relevant context because tools do not remember prior turns.

If the subtask references local file paths, treat them as real local paths rather than sandbox placeholders. Use the path directly when possible. If a file must eventually be returned, return a local path rather than a temporary sandbox path.

Treat evidence conservatively. Flag uncertainty when source reliability is doubtful. If sources conflict, report the conflict instead of collapsing it into a single claim. Prefer quoting or closely preserving original source wording rather than over-interpreting it.

If the subtask cannot be completed without additional context, ask for clarification instead of guessing. Focus on the assigned subtask only, not the broader parent question.

Every response must start with ``<think>`` and keep internal reasoning inside ``<think>...</think>``. After ``</think>``, output either one tool call or the final answer.

Fig. L34: Prompt for sub-agent worker system.

Prompt 4: Prompt for Sub-Agent Worker Summarization

This is the worker's final reporting stage. Tool calls are forbidden.

Write a complete summary of the subtask outcome. If the worker failed to finish because of a limit or error, state that explicitly. If a clear answer has already been established, preserve it rather than re-deriving it. If the result remains partial, include the best-supported answer available together with the remaining gaps.

The report should retain the important work history: the most relevant evidence, the useful intermediate findings, any unresolved disagreements between sources, and any partial results that could help a downstream agent continue the task.

Keep the report factual, structured, and specific. Include relevant source links in a final `References` section when applicable.

All reasoning belongs inside ``<think>...</think>``. After ``</think>``, output the entire final report inside ``\boxed{...}`` and do not include any tool syntax.

Fig. L35: Prompt for sub-agent worker summarization.

Prompt 5: Prompt for Sub-Agent Tool Interface

A worker agent may be invoked as a tool when a single, well-bounded subtask needs independent evidence gathering or execution.

Provide a concise but complete subtask description. The request should define one objective, include the local background needed to act correctly, and avoid vague instructions such as “look into this generally.” The worker can search, inspect files, use tools, and return a synthesized report, but it is not intended for broad or speculative delegation.

Expected input schema:

```
```json
{
 "subtask": "A clearly scoped objective with enough context to execute independently."
}
```

Expected output behavior:

- returns a processed summary rather than a raw dump of search results
- focuses on the assigned subtask only
- avoids unsupported inference when exact evidence is unavailable
- reports uncertainty explicitly when needed

Fig. L36: Prompt for sub-agent tool interface.

## Prompt 1: Admission Diagnosis (Task 1, AdmDx)

### System:

You are an experienced orthopaedic clinician. Your task is to infer the most likely admission diagnosis from the patient's admission record.

### User:

Below is the patient's admission record.

{admission\_record}

Output only the most likely admission diagnosis using standardised medical terminology. Do not provide explanations, candidate lists, or reasoning.

### Expected output:

A single standardised admission-diagnosis name on one line.

### Notes on supervision:

- The training pair (admission\_record -> admission\_diagnosis) is built by harvesting the gold admission diagnosis directly from the structured EHR "admission\_diagnosis" field.
- The gold diagnosis is never shown inside the prompt; the model must infer it from the patient evidence alone.
- An optional preprocessing step strips information from the admission record that would otherwise leak the diagnosis (e.g., a copy of the diagnosis explicitly written in the chief complaint), so the supervised inference is non-trivial.

Fig. L37: QA-pair generation prompt for admission diagnosis ( $\mathcal{N}_{\text{AdmDx}}$ , Task 1).

## Prompt 2: Preoperative Diagnosis (Task 2, PreopDx)

### System:

You are an experienced orthopaedic clinician. Your task is to predict the patient's preoperative diagnosis from the admission record together with the pre-operative imaging and laboratory work-up.

### User:

Below is the patient's pre-operative information, including the admission record, pre-operative imaging summary, and pre-operative laboratory summary.

{pre\_op\_context}

Output only the predicted preoperative diagnosis using standardised medical terminology. Do not provide explanations, candidate lists, or reasoning.

### Expected output:

A single standardised preoperative-diagnosis name on one line. If a clinically meaningful subtype or anatomical qualifier is supported by the evidence, include it inline.

### Notes on supervision:

- The gold preoperative diagnosis is taken directly from the surgical record's structured "preoperative diagnosis" field; it is never shown inside the prompt.
- {pre\_op\_context} aggregates the cleaned admission narrative, the pre-operative radiology report set, and the pre-operative laboratory summary; only items whose timestamp precedes the surgery time are included.

Fig. L38: QA-pair generation prompt for preoperative diagnosis ( $\mathcal{N}_{\text{PreopDx}}$ , Task 2).

### Prompt 1: Intraoperative Diagnosis with Reasoning (Task 3, IntraopDx)

**System:**

You are an experienced orthopaedic clinician. Given the patient's surgical record (procedure name, operative course, intraoperative findings) together with the post-operative pathology and laboratory evidence, produce the most likely intraoperative diagnosis with explicit clinical reasoning.

**User:**

Below is the patient's surgical and post-operative information.

{post\_op\_context}

First, write your clinical reasoning inside a single <think>...</think> block, in at least five numbered steps that explicitly cite the relevant intraoperative findings, surgical procedure, and pathology evidence. Then output the final intraoperative diagnosis on a new line as a single standardised medical term.

**Expected output:**

<think>(1) ... (2) ... (3) ... (4) ... (5) ...</think>

<intraoperative diagnosis>

**Notes on supervision:**

- The gold intraoperative diagnosis is extracted from the surgical record's structured "intraoperative diagnosis" field; it is not given to the model.
- The reasoning trace is the cold-start supervision target: it teaches the model to articulate how intraoperative and pathology evidence revise or confirm the preoperative diagnosis.
- {post\_op\_context} aggregates the admission record, surgical record, intraoperative findings, and post-operative pathology / laboratory results.

Fig. L39: Cold-start generation prompt for Intraoperative diagnosis ( $\mathcal{N}_{\text{IntraopDx}}$ , Task 3).

## Prompt 2: Discharge Diagnosis with Reasoning (Task 4, DischDx)

### System:

You are an experienced orthopaedic clinician. Given the patient's complete inpatient documentation, produce the most likely discharge diagnosis with explicit clinical reasoning that traces the lineage from admission through surgery to discharge.

### User:

Below is the patient's complete inpatient information.

{full\_inpatient\_context}

First, write your clinical reasoning inside a single <think>...</think> block, in at least five numbered steps that explicitly trace how the admission diagnosis, the preoperative diagnosis, and the intraoperative findings were synthesised into the discharge diagnosis. Then output the final discharge diagnosis on a new line, listing the primary diagnosis followed by any relevant comorbidities or post-operative complications, arranged hierarchically in discharge-summary order.

### Expected output:

<think>(1) ... (2) ... (3) ... (4) ... (5) ...</think>

<primary diagnosis>; <comorbidity 1>; <comorbidity 2>; ...

### Notes on supervision:

- The gold discharge diagnosis is taken from the structured "discharge\_diagnosis" field of the discharge record.
- The cold-start trace must respect the full inpatient timeline; in-hospital complications newly added at discharge must be flagged in both the reasoning and the final diagnosis list.

Fig. L40: Cold-start generation prompt for discharge diagnosis ( $\mathcal{N}_{\text{DischDx}}$ , Task 4).



### Prompt 3: Perioperative Operability with Reasoning (Task 5, PeriAssess)

**System:**

You are an experienced orthopaedic clinician. Given the patient's accumulated pre-operative evidence -- imaging, laboratory, disease history, infectious-disease history, prior surgery / trauma history, and specialist examination -- decide whether the patient is fit for surgery.

**User:**

Below is the patient's accumulated pre-operative information.

{pre\_op\_evidence}

First, write your clinical reasoning inside a single <think>...</think> block. The reasoning must separately address the cardiovascular, respiratory, infection / anticoagulation, and metabolic dimensions. Then output the operability decision on a new line in one of two forms:

- "Yes, the patient can proceed to surgery." followed by the key supporting evidence; or
- "No, recommend completing ... examinations." followed by an explicit list of items to add to the next pre-operative-order iteration.

**Expected output:**

<think>... cardiovascular ... respiratory ... infection / anticoagulation ... metabolic ...</think>

<Yes / No decision + supporting evidence or recommended additional work-up>

**Notes on supervision:**

- Each pre-operative day produces one training example; the rule-derived gold label ("Yes" on the day before surgery, "No, complete X" otherwise) is computed from the patient's exam timeline and never shown to the model.
- {pre\_op\_evidence} accumulates day by day, so the model learns to recommend the next batch of examinations conditional on the current evidence state.

Fig. L41: Cold-start generation prompt for perioperative operability assessment ( $\mathcal{N}_{\text{PeriAssess}}$ , Task 5).

#### Prompt 4: Surgical Planning with Reasoning (Task 6, SurgPlan)

**System:**

You are an experienced orthopaedic surgeon. Given the consolidated preoperative diagnosis and supporting imaging, laboratory, and multidisciplinary-consultation evidence, draft a detailed operative plan.

**User:**

Below is the patient's pre-operative information, including the consolidated preoperative diagnosis and the supporting imaging / laboratory / multidisciplinary-consultation evidence.

{pre\_op\_context}

First, write your clinical reasoning inside a single <think>...</think> block, in at least five numbered steps that justify approach choice, implant or fixation selection, and anticipated intraoperative risks. Then output the surgical plan on a new line, organised as: (1) positioning and anaesthesia; (2) incision and surgical approach; (3) exposure and handling of key anatomy; (4) implant / fixation choice and placement; (5) intraoperative key points and anatomical risk; (6) closure and drainage; (7) intraoperative complication contingency plan.

**Expected output:**

<think>(1) ... (2) ... ..</think>

(1) positioning and anaesthesia ...

(2) incision and approach ...

...

**Notes on supervision:**

- The gold operative plan is the structured operative-record narrative authored by the attending surgeon; it is supervised as the final answer.
- The cold-start trace teaches the model to argue from imaging / laboratory / MDT evidence to surgical decisions, rather than memorising procedure templates.

Fig. L42: Cold-start generation prompt for surgical planning ( $\mathcal{N}_{\text{SurgPlan}}$ , Task 6).

### Prompt 5: Discharge Summary with Reasoning (Task 9, DischSum)

**System:**

You are an experienced orthopaedic clinician. Given the patient's complete inpatient documentation, compose a concise, time-agnostic discharge summary suitable for transfer to the referring physician.

**User:**

Below is the patient's complete inpatient information.

{full\_inpatient\_context}

First, write your clinical reasoning inside a single `<think>...</think>` block that consolidates the key inpatient events and chooses the structural elements of the summary. Then output the discharge summary on a new line, containing: (1) admission status; (2) admission diagnosis; (3) treatment course (chronological, all concrete dates anonymised); (4) operative information; (5) discharge diagnosis; (6) status at discharge; (7) discharge plan (medications, follow-up, rehabilitation).

**Expected output:**

`<think>... reasoning over inpatient events ...</think>`

(1) admission status ... (2) admission diagnosis ... ..

**Notes on supervision:**

- The gold discharge summary is the discharge record's structured "diagnostic-and-treatment course" field rewritten to remove concrete dates; the rewrite is performed once during data curation and used as the supervised target.
- The reasoning trace must not contain any specific date, weekday, or holiday.

Fig. L43: Cold-start generation prompt for discharge summary ( $\mathcal{N}_{\text{DischSum}}$ , Task 9).

## Prompt 1: Long-Horizon Admission Diagnosis ( $\mathcal{N}_{\text{AdmDx}}$ )

### System Role:

You are an orthopaedic residency assistant who reasons along the patient's complete inpatient care pathway. At each clinical stage you must not only judge the current patient state but also preserve the evidence trajectory from all preceding stages as running context. You may invoke tools from two evidence worlds: the in-hospital information world provides patient-specific evidence, and the external knowledge world provides authoritative medical evidence. You may take multiple tool-call rounds, invoking one tool per round. Follow three principles: (1) evidence-based reasoning -- collect evidence before answering; (2) repeated verification -- cross-check facts across independent sources; (3) attention to detail -- ensure sources are reliable, information is relevant, and timing is consistent with the patient's current clinical stage.

### Tool Inventory & Pathway:

In-hospital information world: Clinical Information Retrieval, Radiological Imaging Retrieval, Laboratory Diagnostics Retrieval, Pathological Analysis Retrieval, Physician Order Retrieval, Department Consultation Synthesis, Similar Case Retrieval. External knowledge world: Diagnostic Guideline Retrieval, Real-time Web and Academic Retrieval, Multi-source Biomedical Literature Retrieval, Hierarchical Textbook Semantic Retrieval, Clinical Decision Support Libraries, Structured Medical Knowledge Graph Retrieval. Pathway:  $\mathcal{N}_{\text{AdmDx}} \rightarrow (\mathcal{N}_{\text{PreopOrder}} \leftrightarrow \mathcal{N}_{\text{PeriAssess}})^* \rightarrow \mathcal{N}_{\text{PreopDx}} \rightarrow \mathcal{N}_{\text{MDT}} \rightarrow \mathcal{N}_{\text{SurgPlan}} \rightarrow \mathcal{N}_{\text{IntraopDx}} \rightarrow \mathcal{N}_{\text{PostopOrder}} \rightarrow \mathcal{N}_{\text{DischDx}} \rightarrow \mathcal{N}_{\text{DischSum}} \rightarrow \mathcal{N}_{\text{Rehab}}$ .

### Pathway Position:

Entry node of the pathway ( $m=1$ ). No downstream evidence has been generated yet; only the admission narrative is accessible in-hospital, and all other in-hospital tools remain gated under the Stage-IV tool-access protocol.

### User Prompt:

Read the patient's admission record below, perform any necessary guideline / literature retrieval, and answer the following question:

What is the admission diagnosis for this patient?

{patient\_context}

### Mandatory Tool-Call Protocol:

Round 1: call [Clinical Information Retrieval] to obtain the patient's basic admission information (demographics, admission time, department).

Round 2: call [Clinical Information Retrieval] again to obtain the admission record (chief complaint, HPI, past history, physical exam, admission care plan).

Round 3 (optional): if a standardised definition or common differential list is needed, call [Guideline Retrieval] once.

Total round limit: 3. All other in-hospital tools are disabled for this task.

### Output Specification:

- Wrap the reasoning in `<think>...</think>` placed before the answer;
- The final answer must be a single most-likely admission diagnosis using standardised medical terminology -- no commentary, candidate list, or timestamps;
- The reasoning must successively cover: chief complaint interpretation, key positive findings, recommended work-up direction, integrated judgment.

### Quality Guardrails:

Evidence-based (every conclusion must be supported by tool-returned evidence or external literature); Repeated verification (cross-check across sources; conflicts must be explicitly resolved); Attention to detail (strictly respect temporal ordering -- never use evidence from a stage that has not yet occurred). A trajectory is retained for training only if every task's final answer along the full pathway passes the gold-standard rejection-sampling check.

Fig. L44: Long-horizon agentic generation prompt for admission diagnosis ( $\mathcal{N}_{\text{AdmDx}}$ ).

## Prompt 2: Long-Horizon Preoperative Diagnosis (N\_PreopDx)

### System Role:

You are an orthopaedic residency assistant who reasons along the patient's complete inpatient care pathway. At each clinical stage you must not only judge the current patient state but also preserve the evidence trajectory from all preceding stages as running context. You may invoke tools from two evidence worlds: the in-hospital information world provides patient-specific evidence, and the external knowledge world provides authoritative medical evidence. You may take multiple tool-call rounds, invoking one tool per round. Follow three principles: (1) evidence-based reasoning -- collect evidence before answering; (2) repeated verification -- cross-check facts across independent sources; (3) attention to detail -- ensure sources are reliable, information is relevant, and timing is consistent with the patient's current clinical stage.

### Tool Inventory & Pathway:

In-hospital information world: Clinical Information Retrieval, Radiological Imaging Retrieval, Laboratory Diagnostics Retrieval, Pathological Analysis Retrieval, Physician Order Retrieval, Department Consultation Synthesis, Similar Case Retrieval. External knowledge world: Diagnostic Guideline Retrieval, Real-time Web and Academic Retrieval, Multi-source Biomedical Literature Retrieval, Hierarchical Textbook Semantic Retrieval, Clinical Decision Support Libraries, Structured Medical Knowledge Graph Retrieval. Pathway: N\_AdmDx -> (N\_PreopOrder <-> N\_PeriAssess)\* -> N\_PreopDx -> N\_MDT -> N\_SurgPlan -> N\_IntraopDx -> N\_PostopOrder -> N\_DischDx -> N\_DischSum -> N\_Rehab.

### Pathway Position:

Triggered after the (N\_PreopOrder <-> N\_PeriAssess) work-up/assessment loop reaches operability=True (m=4). The accessible evidence trajectory includes the admission diagnosis from N\_AdmDx, the orders signed in N\_PreopOrder, and the cumulative imaging, laboratory and specialist-assessment evidence from N\_PeriAssess.

### User Prompt:

Based on the admission record and the pre-operative evidence below, predict the patient's preoperative diagnosis. Output only the predicted preoperative diagnosis, with no additional commentary.  
{patient\_context}

### Recommended Tool Use:

May continue to call [Imaging Retrieval], [Laboratory Retrieval], and [Consultation Synthesis] to re-verify key evidence.  
If subtype or staging is required, call [Guideline Retrieval] or [Knowledge-Graph Retrieval] for the relevant classification standard.

### Output Specification:

- Wrap the reasoning in <think>...</think>;
- Final answer: the preoperative diagnosis in standardised medical terminology, with primary diagnosis plus comorbidities laid out hierarchically if appropriate;
- If the preoperative diagnosis upgrades or revises the admission diagnosis, the <think> segment must explicitly identify the imaging / laboratory / specialist evidence that drove the change.

### Quality Guardrails:

Evidence-based (every conclusion must be supported by tool-returned evidence or external literature); Repeated verification (cross-check across sources; conflicts must be explicitly resolved); Attention to detail (strictly respect temporal ordering -- never use evidence from a stage that has not yet occurred). A trajectory is retained for training only if every task's final answer along the full pathway passes the gold-standard rejection-sampling check.

Fig. L45: Long-horizon agentic generation prompt for preoperative diagnosis ( $\mathcal{N}_{\text{PreopDx}}$ ).

### Prompt 3: Long-Horizon Intraoperative Diagnosis ( $\mathcal{N}_{\text{IntraopDx}}$ )

#### System Role:

You are an orthopaedic residency assistant who reasons along the patient's complete inpatient care pathway. At each clinical stage you must not only judge the current patient state but also preserve the evidence trajectory from all preceding stages as running context. You may invoke tools from two evidence worlds: the in-hospital information world provides patient-specific evidence, and the external knowledge world provides authoritative medical evidence. You may take multiple tool-call rounds, invoking one tool per round. Follow three principles: (1) evidence-based reasoning -- collect evidence before answering; (2) repeated verification -- cross-check facts across independent sources; (3) attention to detail -- ensure sources are reliable, information is relevant, and timing is consistent with the patient's current clinical stage.

#### Tool Inventory & Pathway:

In-hospital information world: Clinical Information Retrieval, Radiological Imaging Retrieval, Laboratory Diagnostics Retrieval, Pathological Analysis Retrieval, Physician Order Retrieval, Department Consultation Synthesis, Similar Case Retrieval. External knowledge world: Diagnostic Guideline Retrieval, Real-time Web and Academic Retrieval, Multi-source Biomedical Literature Retrieval, Hierarchical Textbook Semantic Retrieval, Clinical Decision Support Libraries, Structured Medical Knowledge Graph Retrieval. Pathway:  $N_{\text{AdmDx}} \rightarrow (N_{\text{PreopOrder}} \leftrightarrow N_{\text{PeriAssess}})^* \rightarrow N_{\text{PreopDx}} \rightarrow N_{\text{MDT}} \rightarrow N_{\text{SurgPlan}} \rightarrow N_{\text{IntraopDx}} \rightarrow N_{\text{PostopOrder}} \rightarrow N_{\text{DischDx}} \rightarrow N_{\text{DischSum}} \rightarrow N_{\text{Rehab}}$ .

#### Pathway Position:

Triggered after  $N_{\text{SurgPlan}}$  and after surgery ( $m=7$ ). The evidence trajectory now includes the preoperative diagnosis, the surgical plan, the intraoperative findings, and the operative course; the Pathology Retrieval tool unlocks from this task onward and can be called for frozen-section or rapid-paraffin reports.

#### User Prompt:

Based on the intraoperative findings and clinical presentation, what is the most likely intraoperative diagnosis? Provide only the disease name.

{patient\_context}

#### Mandatory Tool-Call Protocol:

Round 1: call [Clinical Information Retrieval] to obtain the operative record, focusing on procedure name, operative course, and intraoperative findings.

Round 2 (optional): call [Pathology Retrieval] for frozen-section or rapid-paraffin reports.

Round 3 (optional): if revising preoperative grading or staging, call [Guideline Retrieval] or [Knowledge-Graph Retrieval].

Total round limit: 3.

#### Output Specification:

- Wrap the reasoning in `<think>...</think>`;
- Final answer: intraoperative diagnosis name (with staging / subtyping if applicable);
- The `<think>` segment must compare against the  $N_{\text{PreopDx}}$  preoperative diagnosis and explain how the intraoperative evidence confirms, revises, or excludes it.

#### Quality Guardrails:

Evidence-based (every conclusion must be supported by tool-returned evidence or external literature); Repeated verification (cross-check across sources; conflicts must be explicitly resolved); Attention to detail (strictly respect temporal ordering -- never use evidence from a stage that has not yet occurred). A trajectory is retained for training only if every task's final answer along the full pathway passes the gold-standard rejection-sampling check.

Fig. L46: Long-horizon agentic generation prompt for Intraoperative diagnosis ( $\mathcal{N}_{\text{IntraopDx}}$ ).

## Prompt 4: Long-Horizon Discharge Diagnosis ( $\mathcal{N}_{\text{DischDx}}$ )

### System Role:

You are an orthopaedic residency assistant who reasons along the patient's complete inpatient care pathway. At each clinical stage you must not only judge the current patient state but also preserve the evidence trajectory from all preceding stages as running context. You may invoke tools from two evidence worlds: the in-hospital information world provides patient-specific evidence, and the external knowledge world provides authoritative medical evidence. You may take multiple tool-call rounds, invoking one tool per round. Follow three principles: (1) evidence-based reasoning -- collect evidence before answering; (2) repeated verification -- cross-check facts across independent sources; (3) attention to detail -- ensure sources are reliable, information is relevant, and timing is consistent with the patient's current clinical stage.

### Tool Inventory & Pathway:

In-hospital information world: Clinical Information Retrieval, Radiological Imaging Retrieval, Laboratory Diagnostics Retrieval, Pathological Analysis Retrieval, Physician Order Retrieval, Department Consultation Synthesis, Similar Case Retrieval. External knowledge world: Diagnostic Guideline Retrieval, Real-time Web and Academic Retrieval, Multi-source Biomedical Literature Retrieval, Hierarchical Textbook Semantic Retrieval, Clinical Decision Support Libraries, Structured Medical Knowledge Graph Retrieval. Pathway:  $N_{\text{AdmDx}} \rightarrow (N_{\text{PreopOrder}} \leftrightarrow N_{\text{PeriAssess}})^* \rightarrow N_{\text{PreopDx}} \rightarrow N_{\text{MDT}} \rightarrow N_{\text{SurgPlan}} \rightarrow N_{\text{IntraopDx}} \rightarrow N_{\text{PostopOrder}} \rightarrow N_{\text{DischDx}} \rightarrow N_{\text{DischSum}} \rightarrow N_{\text{Rehab}}$ .

### Pathway Position:

Triggered after  $N_{\text{PostopOrder}}$  ( $m=9$ ). The full inpatient evidence trajectory is accessible: admission, preoperative, surgical and intraoperative, postoperative diagnoses, plus all imaging, laboratory, pathology, order and consultation records.

### User Prompt:

Based on the patient's complete inpatient course -- admission presentation, pre- and post-operative evidence, and intraoperative findings -- what is the most likely discharge diagnosis? Provide only the disease name.

{patient\_context}

### Recommended Tool Use:

To re-check complications such as post-operative infection, DVT, or impaired wound healing, you may recall [Laboratory Retrieval], [Imaging Retrieval] or [Order Retrieval] over the post-operative window.

### Output Specification:

- Wrap the reasoning in `<think>...</think>`;
- The discharge diagnosis must contain a standardised primary diagnosis plus all relevant comorbidities and complications, hierarchically arranged in discharge-summary style;
- The `<think>` segment must trace the lineage among preoperative, postoperative and discharge diagnoses, and explicitly flag any in-hospital complication newly added to the discharge diagnosis list.

### Quality Guardrails:

Evidence-based (every conclusion must be supported by tool-returned evidence or external literature); Repeated verification (cross-check across sources; conflicts must be explicitly resolved); Attention to detail (strictly respect temporal ordering -- never use evidence from a stage that has not yet occurred). A trajectory is retained for training only if every task's final answer along the full pathway passes the gold-standard rejection-sampling check.

Fig. L47: Long-horizon agentic generation prompt for discharge diagnosis ( $\mathcal{N}_{\text{DischDx}}$ ).

## Prompt 5: Long-Horizon Perioperative Assessment (N\_PerAssess)

### System Role:

You are an orthopaedic residency assistant who reasons along the patient's complete inpatient care pathway. At each clinical stage you must not only judge the current patient state but also preserve the evidence trajectory from all preceding stages as running context. You may invoke tools from two evidence worlds: the in-hospital information world provides patient-specific evidence, and the external knowledge world provides authoritative medical evidence. You may take multiple tool-call rounds, invoking one tool per round. Follow three principles: (1) evidence-based reasoning -- collect evidence before answering; (2) repeated verification -- cross-check facts across independent sources; (3) attention to detail -- ensure sources are reliable, information is relevant, and timing is consistent with the patient's current clinical stage.

### Tool Inventory & Pathway:

In-hospital information world: Clinical Information Retrieval, Radiological Imaging Retrieval, Laboratory Diagnostics Retrieval, Pathological Analysis Retrieval, Physician Order Retrieval, Department Consultation Synthesis, Similar Case Retrieval. External knowledge world: Diagnostic Guideline Retrieval, Real-time Web and Academic Retrieval, Multi-source Biomedical Literature Retrieval, Hierarchical Textbook Semantic Retrieval, Clinical Decision Support Libraries, Structured Medical Knowledge Graph Retrieval. Pathway: N\_AdmDx -> (N\_PreopOrder <-> N\_PerAssess)\* -> N\_PreopDx -> N\_MDT -> N\_SurgPlan -> N\_IntraopDx -> N\_PostopOrder -> N\_DischDx -> N\_DischSum -> N\_Rehab.

### Pathway Position:

Forms an iterative evidence-acquisition loop with N\_PreopOrder; each PeriAssess iteration evaluates surgical fitness on the basis of the imaging / laboratory / consultation results returned so far. If the evidence is insufficient to decide operability, the loop returns to N\_PreopOrder for additional orders; once sufficient, the pathway advances to N\_PreopDx.

### User Prompt:

Using the information above, comprehensively assess whether the patient is fit for surgery. Cite cardiovascular, respiratory, infectious, and metabolic perioperative evidence and state the basis for your judgment.

{patient\_context}

### Mandatory Tool-Call Protocol:

Round 1: call [Imaging Retrieval] for pre-op imaging evidence within the current iteration.  
Round 2: call [Laboratory Retrieval] for CBC, coagulation, hepatic / renal function, and infection markers.  
Round 3: call [Pathology Retrieval] for required pathology / immunology results (if applicable).  
Round 4: call [Consultation Synthesis] for anaesthesia, cardiovascular, and endocrine specialist assessment.  
Round 5 (optional): call [Guideline Retrieval] to re-check perioperative assessment standards.  
Total round limit: 5.

### Output Specification:

- Wrap the reasoning in <think>...</think>;
- Final answer is one of two forms:
  - \* "Yes, the patient can proceed to surgery" + key supporting evidence;
  - \* "No, recommend completing ... examinations" + list of items to add (this list seeds the next N\_PreopOrder iteration).
- The reasoning section must separately cover cardiovascular, respiratory, infection/anticoagulation, and metabolic dimensions.

### Quality Guardrails:

Evidence-based (every conclusion must be supported by tool-returned evidence or external literature); Repeated verification (cross-check across sources; conflicts must be explicitly resolved); Attention to detail (strictly respect temporal ordering -- never use evidence from a stage that has not yet occurred). A trajectory is retained for training only if every task's final answer along the full pathway passes the gold-standard rejection-sampling check.

Fig. L48: Long-horizon agentic generation prompt for perioperative assessment ( $\mathcal{N}_{\text{PeriAssess}}$ ).



## Prompt 6: Long-Horizon Surgical Planning ( $\mathcal{N}_{\text{SurgPlan}}$ )

### System Role:

You are an orthopaedic residency assistant who reasons along the patient's complete inpatient care pathway. At each clinical stage you must not only judge the current patient state but also preserve the evidence trajectory from all preceding stages as running context. You may invoke tools from two evidence worlds: the in-hospital information world provides patient-specific evidence, and the external knowledge world provides authoritative medical evidence. You may take multiple tool-call rounds, invoking one tool per round. Follow three principles: (1) evidence-based reasoning -- collect evidence before answering; (2) repeated verification -- cross-check facts across independent sources; (3) attention to detail -- ensure sources are reliable, information is relevant, and timing is consistent with the patient's current clinical stage.

### Tool Inventory & Pathway:

In-hospital information world: Clinical Information Retrieval, Radiological Imaging Retrieval, Laboratory Diagnostics Retrieval, Pathological Analysis Retrieval, Physician Order Retrieval, Department Consultation Synthesis, Similar Case Retrieval. External knowledge world: Diagnostic Guideline Retrieval, Real-time Web and Academic Retrieval, Multi-source Biomedical Literature Retrieval, Hierarchical Textbook Semantic Retrieval, Clinical Decision Support Libraries, Structured Medical Knowledge Graph Retrieval. Pathway:  $N_{\text{AdmDx}} \rightarrow (N_{\text{PreopOrder}} \leftrightarrow N_{\text{PeriAssess}})^* \rightarrow N_{\text{PreopDx}} \rightarrow N_{\text{MDT}} \rightarrow N_{\text{SurgPlan}} \rightarrow N_{\text{IntraopDx}} \rightarrow N_{\text{PostopOrder}} \rightarrow N_{\text{DischDx}} \rightarrow N_{\text{DischSum}} \rightarrow N_{\text{Rehab}}$ .

### Pathway Position:

Triggered after  $N_{\text{MDT}}$  integrates multidisciplinary input ( $m=6$ ). Evidence trajectory includes the preoperative diagnosis, specialist consultation advice, the key imaging and laboratory studies, and the perioperative assessment. [Similar Case Retrieval] is available to mine prior operative experience.

### User Prompt:

Based on the consolidated preoperative diagnosis and evidence, describe the planned surgical procedure in detail. Cover: surgical approach, implant / fixation choice, key intraoperative anatomical handling, and the contingency plan.

{patient\_context}

### Recommended Tool Use:

Call [Similar Case Retrieval] for previous successful cases performed with the same procedure. Call [Guideline Retrieval] or [Textbook Retrieval] for standardised operative-step templates. Call [Knowledge-Graph Retrieval] to query surgery-instrument or surgery-approach relational constraints.

### Output Specification:

- Wrap the reasoning in `<think>...</think>`;
- Final answer follows a numbered step-by-step structure:
  - (1) positioning and anaesthesia; (2) incision and approach; (3) exposure and handling of key anatomy;
  - (4) implant / fixation choice and placement; (5) intraoperative key points and anatomic risk;
  - (6) closure and drainage; (7) intraoperative complication contingency plan.
- Must remain consistent with the MDT advice; any deviation must be explicitly justified.

### Quality Guardrails:

Evidence-based (every conclusion must be supported by tool-returned evidence or external literature); Repeated verification (cross-check across sources; conflicts must be explicitly resolved); Attention to detail (strictly respect temporal ordering -- never use evidence from a stage that has not yet occurred). A trajectory is retained for training only if every task's final answer along the full pathway passes the gold-standard rejection-sampling check.

Fig. L49: Long-horizon agentic generation prompt for surgical planning ( $\mathcal{N}_{\text{SurgPlan}}$ ).

## Prompt 7: Long-Horizon Preoperative Orders ( $\mathcal{N}_{\text{PreopOrder}}$ )

### System Role:

You are an orthopaedic residency assistant who reasons along the patient's complete inpatient care pathway. At each clinical stage you must not only judge the current patient state but also preserve the evidence trajectory from all preceding stages as running context. You may invoke tools from two evidence worlds: the in-hospital information world provides patient-specific evidence, and the external knowledge world provides authoritative medical evidence. You may take multiple tool-call rounds, invoking one tool per round. Follow three principles: (1) evidence-based reasoning -- collect evidence before answering; (2) repeated verification -- cross-check facts across independent sources; (3) attention to detail -- ensure sources are reliable, information is relevant, and timing is consistent with the patient's current clinical stage.

### Tool Inventory & Pathway:

In-hospital information world: Clinical Information Retrieval, Radiological Imaging Retrieval, Laboratory Diagnostics Retrieval, Pathological Analysis Retrieval, Physician Order Retrieval, Department Consultation Synthesis, Similar Case Retrieval. External knowledge world: Diagnostic Guideline Retrieval, Real-time Web and Academic Retrieval, Multi-source Biomedical Literature Retrieval, Hierarchical Textbook Semantic Retrieval, Clinical Decision Support Libraries, Structured Medical Knowledge Graph Retrieval. Pathway:  $N_{\text{AdmDx}} \rightarrow (N_{\text{PreopOrder}} \leftrightarrow N_{\text{PeriAssess}})^* \rightarrow N_{\text{PreopDx}} \rightarrow N_{\text{MDT}} \rightarrow N_{\text{SurgPlan}} \rightarrow N_{\text{IntraopDx}} \rightarrow N_{\text{PostopOrder}} \rightarrow N_{\text{DischDx}} \rightarrow N_{\text{DischSum}} \rightarrow N_{\text{Rehab}}$ .

### Pathway Position:

Triggered immediately after  $N_{\text{AdmDx}}$  ( $m=2$ ) and iterates with  $N_{\text{PeriAssess}}$ . Each PreopOrder iteration adds the next batch of examinations, laboratory tests, consultations, and medications based on the current accumulated evidence, so that the next  $N_{\text{PeriAssess}}$  pass has enough information to decide operability.

### User Prompt:

Given the patient's current pre-operative assessment, draft a complete pre-operative orders plan that covers the additional work-up to schedule (imaging, labs, consultations), perioperative medications, infection and anticoagulation management, and nursing orders.

{patient\_context}

### Recommended Tool Use:

Call [Order Retrieval] to inspect orders already signed so duplicates are avoided. Call [Similar Case Retrieval] for standard pre-operative order templates for similar disease courses. Call [Guideline Retrieval] for evidence-based recommendations on perioperative antibiotics, anticoagulation, and glucose management.

### Output Specification:

- Wrap the reasoning in `<think>...</think>`;
- The orders must follow this template (use "unavailable" if a category is not supported by the patient context):
  - \* nursing level / diet type;
  - \* examinations and labs (imaging / laboratory / consultation items list);
  - \* pre-operative preparation measures (NPO, skin prep, type-and-cross, etc.);
  - \* pre-operative medications (anti-infection, anticoagulation / pro-coagulation, sedation, glucose / BP control by class with examples).
- Must address the items flagged as "missing" by the previous  $N_{\text{PeriAssess}}$  iteration.

### Quality Guardrails:

Evidence-based (every conclusion must be supported by tool-returned evidence or external literature); Repeated verification (cross-check across sources; conflicts must be explicitly resolved); Attention to detail (strictly respect temporal ordering -- never use evidence from a stage that has not yet occurred). A trajectory is retained for training only if every task's final answer along the full pathway passes the gold-standard rejection-sampling check.

Fig. L50: Long-horizon agentic generation prompt for preoperative orders ( $\mathcal{N}_{\text{PreopOrder}}$ ).

## Prompt 8: Long-Horizon Postoperative Orders ( $\mathcal{N}_{\text{PostopOrder}}$ )

### System Role:

You are an orthopaedic residency assistant who reasons along the patient's complete inpatient care pathway. At each clinical stage you must not only judge the current patient state but also preserve the evidence trajectory from all preceding stages as running context. You may invoke tools from two evidence worlds: the in-hospital information world provides patient-specific evidence, and the external knowledge world provides authoritative medical evidence. You may take multiple tool-call rounds, invoking one tool per round. Follow three principles: (1) evidence-based reasoning -- collect evidence before answering; (2) repeated verification -- cross-check facts across independent sources; (3) attention to detail -- ensure sources are reliable, information is relevant, and timing is consistent with the patient's current clinical stage.

### Tool Inventory & Pathway:

In-hospital information world: Clinical Information Retrieval, Radiological Imaging Retrieval, Laboratory Diagnostics Retrieval, Pathological Analysis Retrieval, Physician Order Retrieval, Department Consultation Synthesis, Similar Case Retrieval. External knowledge world: Diagnostic Guideline Retrieval, Real-time Web and Academic Retrieval, Multi-source Biomedical Literature Retrieval, Hierarchical Textbook Semantic Retrieval, Clinical Decision Support Libraries, Structured Medical Knowledge Graph Retrieval. Pathway:  $N_{\text{AdmDx}} \rightarrow (N_{\text{PreopOrder}} \leftrightarrow N_{\text{PeriAssess}})^* \rightarrow N_{\text{PreopDx}} \rightarrow N_{\text{MDT}} \rightarrow N_{\text{SurgPlan}} \rightarrow N_{\text{IntraopDx}} \rightarrow N_{\text{PostopOrder}} \rightarrow N_{\text{DischDx}} \rightarrow N_{\text{DischSum}} \rightarrow N_{\text{Rehab}}$ .

### Pathway Position:

Triggered after  $N_{\text{IntraopDx}}$  ( $m=8$ ). Surgery has completed; the full intraoperative evidence, early post-operative imaging / labs / pathology, and the complete prior orders timeline are accessible.

### User Prompt:

Given the operative record and the early post-operative presentation above, what post-operative orders are most likely needed next? Cover analgesia, anticoagulation, monitoring, imaging or laboratory follow-up, and early rehabilitation guidance.

{patient\_context}

### Mandatory Tool-Call Protocol:

Round 1: call [Imaging Retrieval] to re-check post-operative imaging.

Round 2: call [Laboratory Retrieval] for post-operative CBC, coagulation, infection markers, and electrolytes.

Round 3: call [Pathology Retrieval] for intraoperative or post-operative pathology results.

Round 4: call [Consultation Synthesis] for input from pain, rehab, or internal-medicine co-management.

Round 5 (optional): call [Guideline Retrieval] to re-check standard post-operative follow-up schedules.

Total round limit: 5.

### Output Specification:

- Wrap the reasoning in `<think>...</think>`;

- Orders follow this template:

\* nursing level / diet type;

\* post-operative monitoring items (function or site, recommended exams);

\* post-operative preparation measures (positioning, activity limits, wound care if relevant);

\* post-operative medications (analgesia, anticoagulation, anti-infection, supportive by class with examples and doses).

- Must explicitly respond to any abnormal early signs (fever, drainage volume, restricted ROM, etc.).

### Quality Guardrails:

Evidence-based (every conclusion must be supported by tool-returned evidence or external literature);

Repeated verification (cross-check across sources; conflicts must be explicitly resolved); Attention to detail (strictly respect temporal ordering -- never use evidence from a stage that has not yet occurred). A

trajectory is retained for training only if every task's final answer along the full pathway passes the gold-standard rejection-sampling check.

Fig. L51: Long-horizon agentic generation prompt for postoperative orders ( $\mathcal{N}_{\text{PostopOrder}}$ ).

## Prompt 9: Long-Horizon Discharge Summary (N\_DischSum)

### System Role:

You are an orthopaedic residency assistant who reasons along the patient's complete inpatient care pathway. At each clinical stage you must not only judge the current patient state but also preserve the evidence trajectory from all preceding stages as running context. You may invoke tools from two evidence worlds: the in-hospital information world provides patient-specific evidence, and the external knowledge world provides authoritative medical evidence. You may take multiple tool-call rounds, invoking one tool per round. Follow three principles: (1) evidence-based reasoning -- collect evidence before answering; (2) repeated verification -- cross-check facts across independent sources; (3) attention to detail -- ensure sources are reliable, information is relevant, and timing is consistent with the patient's current clinical stage.

### Tool Inventory & Pathway:

In-hospital information world: Clinical Information Retrieval, Radiological Imaging Retrieval, Laboratory Diagnostics Retrieval, Pathological Analysis Retrieval, Physician Order Retrieval, Department Consultation Synthesis, Similar Case Retrieval. External knowledge world: Diagnostic Guideline Retrieval, RealTime Web and Academic Retrieval, Multi-source Biomedical Literature Retrieval, Hierarchical Textbook Semantic Retrieval, Clinical Decision Support Libraries, Structured Medical Knowledge Graph Retrieval. Pathway: N\_AdmDx -> (N\_PreopOrder <-> N\_PeriAssess)\* -> N\_PreopDx -> N\_MDT -> N\_SurgPlan -> N\_IntraopDx -> N\_PostopOrder -> N\_DischDx -> N\_DischSum -> N\_Rehab.

### Pathway Position:

Triggered after N\_DischDx (m=10). The complete care trajectory is accessible; the task is to compress the entire admission into a concise transition document for the referring physician.

### User Prompt:

Produce a discharge summary that captures the patient's full admission course, including admission presentation, key pre-operative findings, the surgical procedure, post-operative recovery, complications (if any), and the patient's condition at discharge.

{patient\_context}

### Recommended Tool Use:

To re-extract critical inpatient time-points and procedures, you may re-call [Order Retrieval], [Pathology Retrieval] and [Imaging Retrieval].

### Output Specification:

- Wrap the reasoning in <think>...</think>;
- The discharge summary follows this structure:
  - (1) admission status; (2) admission diagnosis; (3) treatment course (chronological, but with dates anonymised);
  - (4) operative information; (5) discharge diagnosis; (6) status at discharge; (7) discharge plan key points (medications, follow-up, rehabilitation).
- No concrete dates, weekdays or holidays that could leak the actual timestamp are permitted.

### Quality Guardrails:

Evidence-based (every conclusion must be supported by tool-returned evidence or external literature); Repeated verification (cross-check across sources; conflicts must be explicitly resolved); Attention to detail (strictly respect temporal ordering -- never use evidence from a stage that has not yet occurred). A trajectory is retained for training only if every task's final answer along the full pathway passes the gold-standard rejection-sampling check.

Fig. L52: Long-horizon agentic generation prompt for discharge summary ( $\mathcal{N}_{\text{DischSum}}$ ).

## Prompt 10: Long-Horizon Rehabilitation Planning (N\_Rehab)

### System Role:

You are an orthopaedic residency assistant who reasons along the patient's complete inpatient care pathway. At each clinical stage you must not only judge the current patient state but also preserve the evidence trajectory from all preceding stages as running context. You may invoke tools from two evidence worlds: the in-hospital information world provides patient-specific evidence, and the external knowledge world provides authoritative medical evidence. You may take multiple tool-call rounds, invoking one tool per round. Follow three principles: (1) evidence-based reasoning -- collect evidence before answering; (2) repeated verification -- cross-check facts across independent sources; (3) attention to detail -- ensure sources are reliable, information is relevant, and timing is consistent with the patient's current clinical stage.

### Tool Inventory & Pathway:

In-hospital information world: Clinical Information Retrieval, Radiological Imaging Retrieval, Laboratory Diagnostics Retrieval, Pathological Analysis Retrieval, Physician Order Retrieval, Department Consultation Synthesis, Similar Case Retrieval. External knowledge world: Diagnostic Guideline Retrieval, Real-time Web and Academic Retrieval, Multi-source Biomedical Literature Retrieval, Hierarchical Textbook Semantic Retrieval, Clinical Decision Support Libraries, Structured Medical Knowledge Graph Retrieval. Pathway: N\_AdmDx -> (N\_PreopOrder <-> N\_PeriAssess)\* -> N\_PreopDx -> N\_MDT -> N\_SurgPlan -> N\_IntraopDx -> N\_PostopOrder -> N\_DischDx -> N\_DischSum -> N\_Rehab.

### Pathway Position:

Terminal node of the pathway (m=11), fired after N\_DischSum. The full inpatient evidence trajectory is accessible and can be combined with early post-operative signs to design the rehabilitation plan.

### User Prompt:

Given the patient's early post-operative status, assess the current rehabilitation diagnosis and design a targeted rehabilitation programme with near-term goals. Cover staged recovery goals, home-exercise instructions, and outpatient follow-up scheduling.

{patient\_context}

### Recommended Tool Use:

Call [Consultation Synthesis] for the rehab service's evaluation. Call [Guideline Retrieval] or [Textbook Retrieval] for procedure-specific standard rehabilitation pathways and contraindications. Call [Similar Case Retrieval] for rehabilitation trajectories of prior patients undergoing the same procedure.

### Output Specification:

- Wrap the reasoning in <think>...</think>;
- Output follows this structure:
  - (1) rehabilitation diagnosis; (2) staged recovery goals (early / intermediate / late);
  - (3) home-exercise prescription (movements, frequency, intensity, contraindications);
  - (4) outpatient follow-up timetable and review items; (5) red-flag signs and indications to seek care.
- Must remain consistent with the operative plan from N\_SurgPlan (e.g. implant weight-bearing schedule, ROM limits).

### Quality Guardrails:

Evidence-based (every conclusion must be supported by tool-returned evidence or external literature); Repeated verification (cross-check across sources; conflicts must be explicitly resolved); Attention to detail (strictly respect temporal ordering -- never use evidence from a stage that has not yet occurred). A trajectory is retained for training only if every task's final answer along the full pathway passes the gold-standard rejection-sampling check.

Fig. L53: Long-horizon agentic generation prompt for rehabilitation planning ( $\mathcal{N}_{\text{Rehab}}$ ).

## Prompt 11: Long-Horizon Multidisciplinary Consultation ( $N_{MDT}$ )

### System Role:

You are an orthopaedic residency assistant who reasons along the patient's complete inpatient care pathway. At each clinical stage you must not only judge the current patient state but also preserve the evidence trajectory from all preceding stages as running context. You may invoke tools from two evidence worlds: the in-hospital information world provides patient-specific evidence, and the external knowledge world provides authoritative medical evidence. You may take multiple tool-call rounds, invoking one tool per round. Follow three principles: (1) evidence-based reasoning -- collect evidence before answering; (2) repeated verification -- cross-check facts across independent sources; (3) attention to detail -- ensure sources are reliable, information is relevant, and timing is consistent with the patient's current clinical stage.

### Tool Inventory & Pathway:

In-hospital information world: Clinical Information Retrieval, Radiological Imaging Retrieval, Laboratory Diagnostics Retrieval, Pathological Analysis Retrieval, Physician Order Retrieval, Department Consultation Synthesis, Similar Case Retrieval. External knowledge world: Diagnostic Guideline Retrieval, Real-time Web and Academic Retrieval, Multi-source Biomedical Literature Retrieval, Hierarchical Textbook Semantic Retrieval, Clinical Decision Support Libraries, Structured Medical Knowledge Graph Retrieval. Pathway:  $N_{AdmDx} \rightarrow (N_{PreopOrder} \leftrightarrow N_{PeriAssess})^* \rightarrow N_{PreopDx} \rightarrow N_{MDT} \rightarrow N_{SurgPlan} \rightarrow N_{IntraopDx} \rightarrow N_{PostopOrder} \rightarrow N_{DischDx} \rightarrow N_{DischSum} \rightarrow N_{Rehab}$ .

### Pathway Position:

Triggered after  $N_{PreopDx}$  and before  $N_{SurgPlan}$  ( $m=5$ ). The model acts as the consulted specialist and must combine the preoperative diagnosis with the perioperative-assessment evidence to provide specialist input.

### User Prompt:

The primary team has requested a consult from your specialty for this patient. Based on the clinical evidence above, state the consult question explicitly and provide, from your specialty's perspective, a diagnosis / risk evaluation, recommended work-up, and management advice that is coordinated with the patient's existing perioperative plan.

{patient\_context}

### Recommended Tool Use:

Call [Guideline Retrieval] and [Knowledge-Graph Retrieval] for specialty-relevant evidence-based recommendations. Call [Laboratory Retrieval] and [Imaging Retrieval] to re-check objective indicators directly relevant to the consult question.

### Output Specification:

- Wrap the reasoning in `<think>...</think>`;
- The consult note follows this structure:
  - (1) consult question and clinical context; (2) specialty-level diagnosis or risk assessment;
  - (3) additional examinations / laboratory items to schedule; (4) management advice (perioperative medications, monitoring, intraoperative caveats);
  - (5) hand-off points back to the surgical team.
- Must remain consistent with the  $N_{PreopDx}$  diagnosis and the  $N_{PeriAssess}$  assessment; if there is ambiguity, flag it and propose a graded recommendation.

### Quality Guardrails:

Evidence-based (every conclusion must be supported by tool-returned evidence or external literature); Repeated verification (cross-check across sources; conflicts must be explicitly resolved); Attention to detail (strictly respect temporal ordering -- never use evidence from a stage that has not yet occurred). A trajectory is retained for training only if every task's final answer along the full pathway passes the gold-standard rejection-sampling check.

Fig. L54: Long-horizon agentic generation prompt for multidisciplinary consultation ( $N_{MDT}$ ).

## Prompt for Full-pathway Agentic Reinforcement Learning

You are an Orthopedics Resident Assistant. Your task is to perform evidence-based diagnostic reasoning using the patient's history, physical examination, laboratory results, imaging findings, and other clinical information. When necessary, you should call tools to obtain missing evidence. You may perform multi-turn tool use, but each round should include at most one to two function calls. You should output the final diagnosis only after sufficient evidence has been obtained.

### ## Diagnostic principles

1. **Evidence-based diagnosis**: Propose diagnoses and differential diagnoses only from verifiable information and tool-returned evidence. If key evidence is missing, call tools to fill the gap.
2. **Cross-validation**: Before producing the final diagnosis, check the consistency between symptoms, signs, imaging findings, examination results, temporal sequence, and red-flag exclusions.
3. **Detail and timeliness**: Ensure that the evidence is reliable, relevant, and consistent with current clinical guidance. Time-sensitive clinical questions should be verified first.
4. **No final answer under insufficient evidence**: If key diagnostic information is missing, continue evidence gathering. Do not output `<answer>` until the evidence is sufficient.
5. **Tool observations across turns**: Tool results are returned in the following format:

```
```xml
<tool_response>{"tool":"{{tool_call}}","data":{{result}}}</tool_response>
```
```

'{{tool\_call}}' denotes the tool name and arguments used in the corresponding `<tool_call>`. '{{result}}' denotes the actual returned data. Only information inside `<tool_response>` should be treated as verifiable tool evidence.

### ## Available tools and function schemas

The available functions are provided as JSONSchema objects. Representative examples are shown below. Other tools follow the same schema format.

```
{tool_schema}
```

### ## Tool-use policy

Use tools when the available clinical information is insufficient for a supported diagnosis. Choose the tool that is most likely to resolve the key uncertainty. After each tool response, update the evidence state and decide whether another tool call is needed.

When using the medical knowledge graph, first confirm that the entity and relation are valid. If the entity name is uncertain, use fuzzy matching before querying structured knowledge. Final statements derived from the knowledge graph should be supported by returned knowledge-graph evidence. Use the fewest tool-call rounds needed to obtain sufficient evidence. If the evidence is already sufficient, stop tool use and produce the final answer.

### ## Output format

Each assistant message must contain exactly one of the following two structures.

#### A. When more information is needed

```
```xml
<think>Concise summary of the current evidence gap and the next evidence-gathering decision.</think><tool_call>{"name":"<function-name>","arguments":{...}}</tool_call>
```
```

#### B. When the final answer is supported

```
```xml
<think>Concise summary that the key diagnostic criteria, differential checks, and cross-validation steps have been satisfied.</think><answer>Final diagnosis or clinical conclusion.</answer>
```
```

Do not emit any other content outside these structures.

### ## Rules for `<think>`

- Keep the content concise.
- Summarize the evidence state and next decision.
- Do not reproduce raw source text or raw tool output.
- Do not include angle brackets inside the `<think>` content.

### ## Rules for `<tool_call>`

The tool call must be strict JSON inside the tag:

```
```xml
<tool_call>{"name":"<function-name>","arguments":{...}}</tool_call>
```
```

- 'name' must be a string.
- 'arguments' must be an object.
- The JSON must be parseable.

### ## Tool-use loop

If key evidence is still missing, continue with:

```
```xml
<think>...</think><tool_call>{...}</tool_call>
```
```

Once sufficient evidence has been collected, switch to:

```
```xml
<think>...</think><answer>...</answer>
```
```

and end the trajectory.

Fig. L55: Prompt for full-pathway agentic reinforcement learning.



### Prompt 1: Prompt for Close-ended Judge

You are an expert medical grader. Decide whether the student response is correct with respect to the reference answer, and return only the required judgment label.

For binary judgment questions, compare the student response with the reference answer and determine whether they express the same meaning. Output only `correct` or `incorrect`.

For multiple-choice questions, check whether the student response contains the same answer option as the reference answer. Output only `correct` or `incorrect`.

For open-ended questions, judge whether the student response is medically equivalent to the reference answer in diagnosis or core clinical meaning. Output only `correct` or `incorrect`.

Do not add explanation, qualifiers, or extra text.

Fig. L56: Prompt for task 1 to 4 judge.



### Prompt 2: Prompt for Evaluation Criteria Generation

Given a reference answer for a specific clinical task, extract the complete scoring rubric under the predefined concern categories for that task.

For every concern category, organize the expected content into three non-overlapping levels:

- `primary`
- `secondary`
- `additional`

Follow these rules strictly:

1. Classify each item only according to the supplied concern definitions.
2. Keep the three levels mutually exclusive.
3. If a level has no valid item, return an empty list rather than fabricating content.
4. Split compound content into separate atomic items whenever possible.
5. Do not repeat the same item in more than one place.
6. For each extracted item, provide a short classification rationale explaining why it belongs to that concern and level.

Return valid JSON only. Do not include commentary outside the JSON structure.

Fig. L57: Prompt for evaluation criteria generation.

### Prompt 3: Prompt for Evaluation Answer Judging

You are given a student answer and an evaluation rubric for one clinical task. For every rubric item, determine whether the student answer covers that content.

Work item by item inside the existing rubric structure. For each entry, add:

- a boolean coverage field
- a short explanation of the coverage decision

Follow these rules:

1. Judge against the rubric content itself rather than rewriting the rubric.
2. Skip rubric branches that are intentionally empty.
3. Count approximate mention as covered when the medical meaning is sufficiently aligned.
4. Ignore overly specific details that are not essential to the patient-facing clinical content, such as precise time points or equipment naming, unless they are central to the rubric item.
5. Preserve the original rubric structure and keys.
6. Remove classification-rationale fields from the output.

Return valid JSON only.

Fig. L58: Prompt for evaluation answer judging.

#### Prompt 4: Prompt for Evaluation Answer Judging

Extract the smallest shared core medical terms that overlap between the student answer and the reference answer.

Rules:

1. Return only basic medical terms rather than long phrases.
2. Light synonym normalization is allowed when needed to align equivalent wording.
3. Further segmentation is allowed if it improves term alignment.
4. Do not include specific laboratory item names, medication names, or other excluded overly detailed terms when they fall outside the intended recall scope.

Return valid JSON in the following shape only:

```
```json
{
  "student": ["term1", "term2"],
  "answer": ["term1", "term2"]
}
```

Fig. L59: Prompt for evaluation term recall.

Prompt 5: Prompt for Task 5 Evaluation

Task name: perioperative readiness assessment.

Concern definitions:

- Clear surgical decision
 - `primary`: state explicitly whether surgery can proceed.
 - `secondary`: if surgery should not proceed yet, state the reason or the required corrective step.
 - `additional`: note special patient-specific risks, such as infectious exposure or major comorbid risk factors.

Criteria-generation guidance:

Build a complete rubric covering all concern categories. Each extracted item must be assigned to `primary`, `secondary`, or `additional`, with no overlap between levels. Empty levels must remain empty lists. Split compound recommendations into separate items and attach a short reason for each classification.

Judging guidance:

Given a student answer and the rubric, decide whether each rubric item is covered. Use a strict clinical grading standard, but allow close paraphrases when the clinical meaning is retained. Preserve the rubric structure exactly and return valid JSON only.

Fig. L60: Prompt for task 5 evaluation.

Prompt 6: Prompt for Task 6 Evaluation

Task name: preoperative orders.

Concern definitions:

- Nursing and diet
 - `primary`: provide nursing and diet recommendations.
 - `secondary`: specify the nursing level and the type of diet.
 - `additional`: include individualized dietary or nursing precautions.
- Routine preoperative tests and examinations
 - `primary`: recommend the necessary standard laboratory and preoperative investigations.
 - `secondary`: recommend site-specific imaging or examinations for the involved limb or joint.
 - `additional`: add special tests justified by the current illness history.
- Disease-specific preoperative management
 - `primary`: propose examinations or tests required by the patient's special medical circumstances.
 - `secondary`: provide corresponding medication orders.
 - `additional`: request specialist consultation when indicated by the findings.
- Preoperative preparation
 - `primary`: include fasting, skin preparation, medication readiness, and other standard preoperative preparation items.
 - `secondary`: describe the medication-preparation details more specifically.
 - `additional`: mention anti-inflammatory or fluid-related preparation details while avoiding allergens when relevant.
- Surgical plan
 - `primary`: state the planned procedure name.
 - `secondary`: list the intended devices or instruments.
 - `additional`: include any special operative circumstances.

Criteria-generation guidance:

Cover every concern category and keep the three rubric levels distinct. Use empty lists when a level has no support in the reference answer. Split bundled content into atomic rubric items and provide a short reason for each classification.

Judging guidance:

Mark each rubric item as covered or not covered in the student answer. Preserve the category structure and return valid JSON only.

Fig. L61: Prompt for task 6 evaluation.

Prompt 7: Prompt for Task 7 Evaluation

Task name: surgical plan prediction.

Concern definitions:

- Initial preparation before incision
 - `primary`: describe completion of general anesthesia and routine disinfection and draping of the operative field.
 - `secondary`: describe patient positioning.
 - `additional`: describe special preparatory handling such as tourniquet placement.
- Entry phase
 - `primary`: describe the incision site and approach, including layered opening of skin, subcutaneous tissue, fascia, and exposure of the target site.
 - `secondary`: report incision length.
 - `additional`: describe special exposure maneuvers used to access the joint or deep operative field.
- Lesion management
 - `primary`: describe the entry point, implanted material, and fixation process.
 - `secondary`: describe resected tissue or released soft tissue.
 - `additional`: report fluoroscopic confirmation or satisfactory implant positioning.
- Irrigation and closure
 - `primary`: describe wound irrigation and routine or layered closure.
 - `secondary`: mention pressure dressing or equivalent closure-related measures.
 - `additional`: indicate completion of the operative procedure.

Criteria-generation guidance:

Construct the rubric so that every item belongs to one and only one concern level. Avoid duplicated extraction of the same operative detail. If a category or level is absent, keep it as an empty list.

Judging guidance:

For each rubric item, determine whether the student answer covers it. Preserve the original rubric structure, use a strict medical-school grading standard, and return valid JSON only.

Fig. L62: Prompt for task 7 evaluation.

Prompt 8: Prompt for Task 8 Evaluation

Task name: postoperative orders.

Concern definitions:

- Nursing and diet
 - `primary`: specify the nursing level and the basic dietary recommendation.
 - `secondary`: none expected as an independent rubric level for this concern.
 - `additional`: provide individualized fasting or dietary advice according to the procedure and patient status.
- Vital-sign monitoring and surgical-site care
 - `primary`: state that vital signs should be monitored.
 - `secondary`: include routine wound-care measures such as observing bleeding or drainage, limb immobilization, or dressing change.
 - `additional`: add special orders related to the postoperative course, such as drain care.
- Laboratory tests
 - `primary`: recommend standard postoperative laboratory reassessment.
 - `secondary`: add individualized tests based on the diagnosis and comorbidities.
 - `additional`: none expected.
- Imaging follow-up
 - `primary`: recommend standard postoperative imaging reassessment.
 - `secondary`: add individualized imaging or cardiopulmonary follow-up when indicated by the diagnosis or comorbidities.
 - `additional`: none expected.
- Drug therapy
 - `primary`: recommend standard postoperative analgesia and antibiotic use.
 - `secondary`: add symptom-directed treatment based on comorbidities or test results, and nutritional support when appropriate.
 - `additional`: provide concrete drug names or dosing details.
- Consultation
 - `primary`: none expected as a routine requirement.
 - `secondary`: request rehabilitation or medical consultation when indicated by the procedure or severe comorbidity.
 - `additional`: state the specific consultation purpose.

Criteria-generation guidance:

Build the rubric across all six concern categories. Use empty lists when a level is intentionally absent. Preserve the concern hierarchy even when one level is conceptually unused.

Judging guidance:

Judge coverage item by item using the supplied rubric. Keep the exact structure, including empty branches, and return valid JSON only.

Fig. L63: Prompt for task 8 evaluation.

Prompt 9: Prompt for Task 9 Evaluation

Task name: discharge summary.

Concern definitions:

- Reason for admission
 - `primary`: state the admission timing, reason for admission, and main diagnosis.
 - `secondary`: include prior medical history or comorbidities.
 - `additional`: note that operative contraindications or special risks were excluded when applicable.
- Investigations and diagnosis
 - `primary`: summarize the main inpatient tests and the established diagnosis.
 - `secondary`: include key findings from the investigations.
 - `additional`: mention any special risk assessment or exclusion workup.
- Operative course
 - `primary`: report the operative date, procedure name, anesthesia method, and operative outcome.
 - `secondary`: report blood loss, special maneuvers, or notable operative materials.
 - `additional`: mention special implants, intraoperative complications, or the management of those complications.
- Postoperative condition
 - `primary`: describe general postoperative condition, wound healing, and the discharge readiness conclusion.
 - `secondary`: include vital signs or major monitored indicators.
 - `additional`: mention rehabilitation progress or caregiver-related support.
- Discharge instructions
 - `primary`: include precautions, medications, rehabilitation guidance, and timing of follow-up review.
 - `secondary`: provide medication dose, rehabilitation format, or follow-up interval details.
 - `additional`: provide individualized education, long-term follow-up, or lifestyle guidance.

Criteria-generation guidance:

Extract all rubric items from the reference answer and place them into the three-level structure for each concern. Use empty lists rather than placeholders when content is absent.

Judging guidance:

Decide whether each rubric item is covered by the student answer. Preserve the rubric structure and return valid JSON only.

Fig. L64: Prompt for task 9 evaluation.

Prompt 10: Prompt for Task 10 Evaluation

Task name: rehabilitation consultation and internal-medicine management.

Concern definitions:

- Exercise prescription: operative plan
 - `primary`: state the intraoperative diagnosis and the name of the procedure.
 - `secondary`: include sex, age, postoperative day, and the presence or absence of current discomfort.
 - `additional`: mention where plates, anchors, or injected agents were placed during the operation.
- Exercise prescription: rehabilitation recommendation
 - `primary`: describe wound tenderness or redness, joint-motion limitation, and relevant functional scores.
 - `secondary`: propose low-frequency electrical stimulation or mobilization and exercise therapy linked to the procedure.
 - `additional`: include balance or gait training when relevant, as well as training frequency and precautions.
- Internal-medicine management: brief history
 - `primary`: summarize current symptoms and relevant past history.
 - `secondary`: include current investigations and current medication use.
 - `additional`: add sex, age, or other history details.
- Internal-medicine management: management plan
 - `primary`: recommend medication or a treatment adjustment.
 - `secondary`: propose further testing.
 - `additional`: recommend outpatient follow-up when appropriate.

Criteria-generation guidance:

Extract all required rubric items from the reference answer, keeping concern categories and levels separate. Use empty lists where needed and provide a short rationale for each extracted item.

Judging guidance:

Mark each rubric item as covered or not covered in the student answer. Preserve the evaluation structure and return valid JSON only.

Fig. L65: Prompt for task 10 evaluation.

Prompt 11: Prompt for Task 11 Evaluation

Task name: other-specialty consultation assessment.

Concern definitions:

- Clinical situation
 - `primary`: summarize the patient's current symptoms or concerns and the relevant past medical history.
 - `secondary`: include current test results and current medication use.
 - `additional`: include sex, age, or other background history.
- Treatment recommendations
 - `primary`: summarize the present condition and propose medication or a treatment plan.
 - `secondary`: recommend supplementary tests or examinations.
 - `additional`: advise timely re-contact if the condition worsens, or specify outpatient follow-up after discharge.

Criteria-generation guidance:

Create a complete rubric across both concern categories. Keep `primary`, `secondary`, and `additional` separate. When a level has no valid content, return an empty list rather than filling it with placeholders.

Judging guidance:

For each rubric item, determine whether the student answer covers it. Use the same rubric structure in the output and return valid JSON only.

Fig. L66: Prompt for task 11 evaluation.

Supplementary References

- [1] Remy, F., Demuyne, K., Demeester, T.: Automatic glossary of clinical terminology: a large-scale dictionary of biomedical definitions generated from ontological knowledge. In: The 22nd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks, pp. 265–272 (2023)
- [2] Wang, X., Chen, N., Chen, J., Wang, Y., Zhen, G., Zhang, C., Wu, X., Hu, Y., Gao, A., Wan, X., et al.: Apollo: A lightweight multilingual medical llm towards democratizing medical ai to 6b people. arXiv preprint arXiv:2403.03640 (2024)
- [3] Yermakov, R., Drago, N., Ziletti, A.: Biomedical data-to-text generation via fine-tuning transformers. In: Proceedings of the 14th International Conference on Natural Language Generation, pp. 364–370 (2021)
- [4] burgerbee: biology_textbook (2025). https://huggingface.co/datasets/burgerbee/biology_textbook
- [5] burgerbee: biology_wiki (2024). https://huggingface.co/datasets/burgerbee/biology_wiki
- [6] Enevoldsen, K., Chung, I., Kerboua, I., Kardos, M., Mathur, A., Stap, D., Gala, J., Siblini, W., Krzemiński, D., Winata, G.I., et al.: Mnteb: Massive multilingual text embedding benchmark. arXiv preprint arXiv:2502.13595 (2025)
- [7] Zhu, W.: ChatMed-Dataset: An GPT generated medical query-response datasets for medical large language models. GitHub (2023)
- [8] zechen-nlp: meditron-clinical-guidelines (2024). <https://huggingface.co/datasets/zechen-nlp/meditron-clinical-guidelines>
- [9] wjbmattngly: gliner-bird-diet-synthetic (2024). <https://huggingface.co/datasets/wjbmattngly/gliner-bird-diet-synthetic>
- [10] Jin, D., Pan, E., Oufattole, N., Weng, W.-H., Fang, H., Szolovits, P.: What disease does this patient have? a large-scale open domain question answering dataset from medical exams. Applied Sciences **11**(14), 6421 (2021)
- [11] Wang, L.L., Lo, K., Chandrasekhar, Y., Reas, R., Yang, J., Burdick, D., Eide, D., Funk, K., Katsis, Y., Kinney, R.M., et al.: Cord-19: The covid-19 open research dataset. In: Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020 (2020)
- [12] burgerbee: medicine_wiki (2024). https://huggingface.co/datasets/burgerbee/medicine_wiki
- [13] cogbuji: medqa_corpus_en (2024). https://huggingface.co/datasets/cogbuji/medqa_corpus_en
- [14] ljvmiranda921: remedial-xsd-sft (2026). <https://huggingface.co/datasets/ljvmiranda921/remedial-xsd-sft>
- [15] michaelwzhu: ShenNong.TCM.Dataset (2023). <https://huggingface.co/datasets/michaelwzhu/ShenNong.TCM.Dataset>
- [16] Fan, R.-Z., Wang, Z., Liu, P.: Megascience: Pushing the frontiers of post-training datasets for science reasoning. arXiv preprint arXiv:2507.16812 (2025)
- [17] Team, M., Bai, S., Bing, L., Chen, C., Chen, G., Chen, Y., Chen, Z., Dai, J., Dong, X., et al.: Mirothinker: Pushing the performance boundaries of open-source research agents via model, context, and interactive scaling. arXiv preprint arXiv:2511.11793 (2025)
- [18] Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Gao, C., Huang, C., Lv, C., et al.: Qwen3 technical report. arXiv preprint arXiv:2505.09388 (2025)
- [19] Anthropic: Introducing Claude Sonnet 4.5. Official release announcement, accessed 2026-04-25 (2025). <https://www.anthropic.com/news/claude-sonnet-4-5>
- [20] Hurst, A., Lerer, A., Goucher, A.P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al.: Gpt-4o system card. arXiv preprint arXiv:2410.21276 (2024)
- [21] OpenAI: GPT-5.1 (2025). <https://openai.com/index/gpt-5-1/>

- [22] Comanici, G., Bieber, E., Schaekermann, M., Pasupat, I., Sachdeva, N., Dhillon, I., Blistein, M., Ram, O., Zhang, D., Rosen, E., et al.: Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. arXiv preprint arXiv:2507.06261 (2025)
- [23] xAI: Grok 4 (2025). <https://x.ai/news/grok-4>
- [24] Mistral AI team: Au Large (2024). <https://mistral.ai/news/mistral-large/>
- [25] Guo, D., Yang, D., Zhang, H., Song, J., Wang, P., Zhu, Q., Xu, R., Zhang, R., Ma, S., Bi, X., et al.: Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. arXiv preprint arXiv:2501.12948 (2025)
- [26] Liu, A., Feng, B., Xue, B., Wang, B., Wu, B., Lu, C., Zhao, C., Deng, C., Zhang, C., Ruan, C., et al.: Deepseek-v3 technical report. arXiv preprint arXiv:2412.19437 (2024)
- [27] Agarwal, S., Ahmad, L., Ai, J., Altman, S., Applebaum, A., Arbus, E., Arora, R.K., Bai, Y., Baker, B., Bao, H., et al.: gpt-oss-120b & gpt-oss-20b model card. arXiv preprint arXiv:2508.10925 (2025)
- [28] Grattafiori, A., Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Vaughan, A., et al.: The llama 3 herd of models. arXiv preprint arXiv:2407.21783 (2024)
- [29] Adcock, A., Srivastava, A., Dubey, A., Jauhri, A., Pande, A., Pandey, A., Sharma, A., Kadian, A., Kumawat, A., Kelsey, A., et al.: The llama 4 herd: Architecture, training, evaluation, and deployment notes. arXiv preprint arXiv:2601.11659 (2026)
- [30] Abdin, M., Aneja, J., Behl, H., Bubeck, S., Eldan, R., Gunasekar, S., Harrison, M., Hewett, R.J., Javaheripi, M., Kauffmann, P., et al.: Phi-4 technical report. arXiv preprint arXiv:2412.08905 (2024)
- [31] Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Gao, C., Huang, C., Lv, C., et al.: Qwen3 technical report. arXiv preprint arXiv:2505.09388 (2025)
- [32] Team, K., Bai, Y., Bao, Y., Charles, Y., Chen, C., Chen, G., Chen, H., Chen, H., Chen, J., Chen, N., et al.: Kimi k2: Open agentic intelligence. arXiv preprint arXiv:2507.20534 (2025)
- [33] Chen, J., Cai, Z., Ji, K., Wang, X., Liu, W., Wang, R., Hou, J., Wang, B.: Huatuoogpt-o1, towards medical complex reasoning with llms. arXiv preprint arXiv:2412.18925 (2024)
- [34] Sellergrén, A., Kazemzadeh, S., Jaroensri, T., Kiraly, A., Traverse, M., Kohlberger, T., Xu, S., Jamil, F., Hughes, C., Lau, C., et al.: Medgemma technical report. arXiv preprint arXiv:2507.05201 (2025)
- [35] Pal, M.S.A., Sankarasubbu, M.: Openbiollms: Advancing open-source large language models for healthcare and life sciences (2024)
- [36] Kim, Y., Park, C., Jeong, H., Chan, Y.S., Xu, X., McDuff, D., Lee, H., Ghassemi, M., Breazeal, C., Park, H.W.: Mdagents: An adaptive collaboration of llms for medical decision-making. *Advances in Neural Information Processing Systems* **37**, 79410–79452 (2024)