



GOAL 

June 2026

NEWSLETTER |

Hong Kong College of Emergency Medicine

30th Anniversary "Dr. Matthew TSUI Challenge Cup" held on 12 April 2026





Celebrating 30 Years of Excellence: Inter-Collegial Soccer Tournament Highlights.





JCM APR 2026


Case presentation - AHNH
Dr. LEONG Sek Fong



Guest Speaker
Dr. YAN Wing Wa

OSCE - PMH

Dr. CHAN Ngo Tin James



Date: 1 Apr 2026 (Wed)

Time: 5:30- 7:30pm

Venue: Seminar Room, G/F, Block A, Queen Elizabeth Hospital

Program

5:30pm – 6:15pm Case Presentation by AED, AHNH

6:15pm – 7:00pm “CO poisoning – Role of HBOT centre in mass casualties”


Dr YAN Wing Wa, Director of Hyperbaric Oxygen Therapy Centre, PYNEH

7:05pm – 7:30pm Discussion on OSCE by AED, PMH



JCM MAY 2026



Sponsored lecture
Dr. Martin LEUNG



Case presentation - TMH
Dr. LING Gee Yan

Guest Speaker

Dr. Dorothy CHAN



OSCE - POH

Dr. Ricky KAM

Date: 6 May 2026 (Wed)

Time: 5:30- 7:30pm

Venue: H702, 7/F Block H, Princess Margaret Hospital

Program

5:15pm – 5:30pm “Sharing on management of hyperkalaemia”
(Sponsored lecture by AstraZeneca)

5:30pm – 6:15pm Case Presentation by AED, TMH

6:15pm – 7:00pm “Mandatory Reporting on Child Abuse”

Dr Dorothy CHAN, Associate Consultant,
Department of Paediatrics, PWH

7:05pm – 7:30pm Discussion on OSCE by AED, POH






June 2026

NEWSLETTER

Hong Kong College of
Emergency
Medicine

Event

Free Paper Presentation on 15 May 2026



13 Participants



June 2026

NEWSLETTER |

Hong Kong College
of Emergency
Medicine

Digital EM subcommittee: Agentic AI and OpenClaw

Women's Fellows Chapter Celebrating 100 Women EM Fellows 40 years in EM

When I first arrived in 1986, Accident & Emergency Room was called "casualty". I started at the United Christian Hospital, which was a tiny sub vented mission hospital then, with less resources but visionary, pioneered in CAPD, CNS, bereavement care in AED, a hospital without wall. Missionaries like the late Dr. Peterson, Dr. Russel Clark, Dr. John Hereward, and others laid the foundation of the UCH "spirit", to love with Christ, to serve with excellence.

Casualty was initially under Medicine, Dr Russel Clark. Dr Matthew Kwa was the first consultant of UCH Casualty. Every Saturday, all the medical officers gathered for a lively meeting chaired by Dr. Kwa-always accompanied by steaming Dim Sum. While we audited cases, Dr. Hereward admirably ran the casualty solo. Sharing meals became the heart and soul of our department's camaraderie.

In the 1980s, medical officers certified the deceased on board the ambulance -which was scary; the most terrifying were a man jumped from height and the man whose leg was taken by a shark in the waters in Sai Kung. That practice changed after a certified "jerked" at the mortuary.

Cases presented in Casualty reflected Hong Kong's pulse: sewing-needle injuries, kitchen-knife wounds, and last-minute deliveries. In 2003, SARS arrived-an invisible enemy that claimed colleagues'lives. Fear filled the air, but so did an unwavering spirit. We stood together on the front lines. Amoy Garden Block E was the iconic memory. Two decades later, COVID-19 tests our endurance anew, with overcrowding and historic makeshift beds.

Looking back, four decades in emergency medicine have brought pride, joy, tears, and unbreakable bonds. In the mid-90s, emergency medicine emerged as a specialty. I still remember the stress of preparing for the first FHKCEM exams, coupled with the pregnancy and rotating to orthopaedics. Life as a working female and mother, was tough. Today's wellness programs and greater gender balance — championed by Dr. Y.F. Ho's women fellows chapter-promise a brighter, more fulfilling future. Here's to the next chapter: stronger together, ever ready to serve with compassion and resilience.

Dr TING See Mai



Agentic AI and OpenClaw

April 2026

Prompter and editor: Dr KL Cheung

Disclaimer: The following article, including the graphics and citations, is generated by Hermes Agent running on GPT-5 large language model. The literature search, review, verification and summarisation were conducted by the agent autonomously.

The user is responsible for the prompt specifying the context, the task, the writing style and content outline. The article created by the agent is reviewed and edited by the prompter to ensure its validity, tone and accuracy before publication. Other AI tools were used in the proofreading of the article.

From fancy chatbots to smart co-workers

Artificial intelligence is already embedded in daily life. Many doctors use chatbots to summarise research papers, draft emails, and explain clinical guidelines. Other AI-enabled applications and workflow or automation software had streamlined some the fragmented digital tasks of a busy doctor. These reactive chatbots and applications cannot leave their own sandboxes and required an operator. Agent AI is the latest emerging concept in the AI world, and it extends way beyond just answering questions. While chatbots respond to prompts, agents can act autonomously and continuously in the background, operating within boundaries defined by the user. A sophisticated agent can design, refine and execute the workflow with the tools it has decided to use, based on objectives of the task, the given context and memory. It orchestrates the acquisition, extraction, and integration of information across multiple services and interfaces simultaneously. At the same time, it must be vigilant to errors, prompting human oversight when needed [1,2].

Figure 1. Chatbot, workflow, agent

A plain-language view of why agentic AI feels different from ordinary chat.

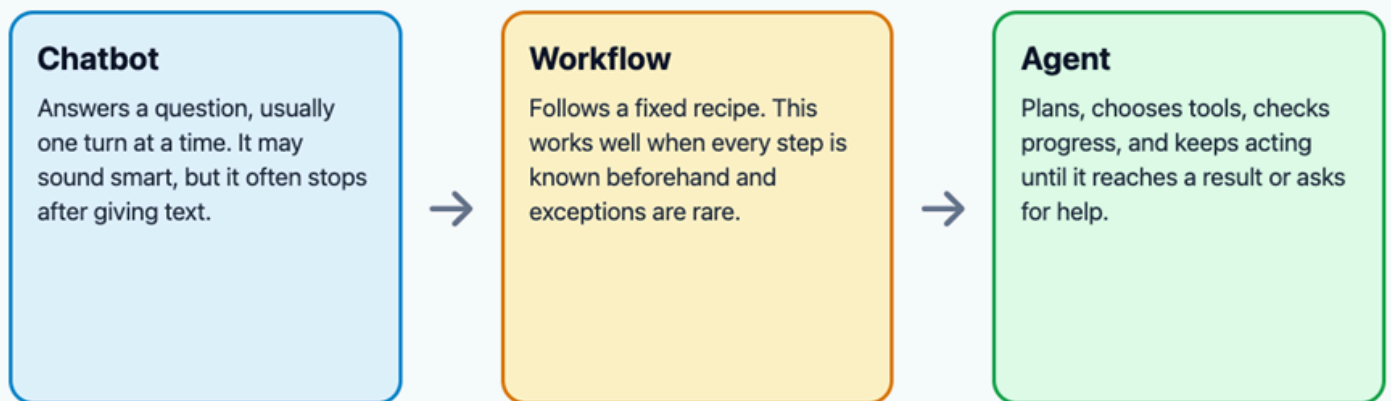


Figure 1. Evolution from AI chatbot, AI workflow, and AI agent. Original figure created for this article. Concept source pages: OpenAI practical guide to building agents (<https://openai.com/business/guides-and-resources/a-practical-guide-to-building-ai-agents/>) and Anthropic Building Effective AI Agents (<https://www.anthropic.com/research/building-effective-agents>).



Agentic AI transforms large language models from passive tools into active assistants. Advanced AI agents can create output with quality approaching professional standards. Some systems are even designed to adapt to and remember user preferences through interactions. By retaining information about schedules, habits, and preferences, they can operate in contextually relevant ways and remain aligned with user expectations, like an intelligent and adaptive, yet virtual and tireless co-worker. Looking ahead, these capabilities may extend further into professional domains, including the development of clinical AI agents to support healthcare practice.

The Lobster Hype

The open-source agent framework OpenClaw is a personal AI assistant running locally on a user’s machine and operates through messaging platforms such as WhatsApp and Telegram, while maintaining native connections to applications on the host device for agentic operations [3,4]. Users can “text” their own machine for instructions, and it would respond and update about the assignment -- like a digital butler of sorts, sometimes likened to “keeping a lobster pet.” The hype it created has led to a surge in demand for small-form personal computers. While OpenClaw appeals to users who seek fine-grained control and are comfortable running systems locally, other options are available. As a competition to OpenClaw, the Hermes Agent promises better long-term memory and self-improving automations [5].

Figure 3. OpenClaw as a gateway agent

One assistant can sit across several chat surfaces, models, and tools.

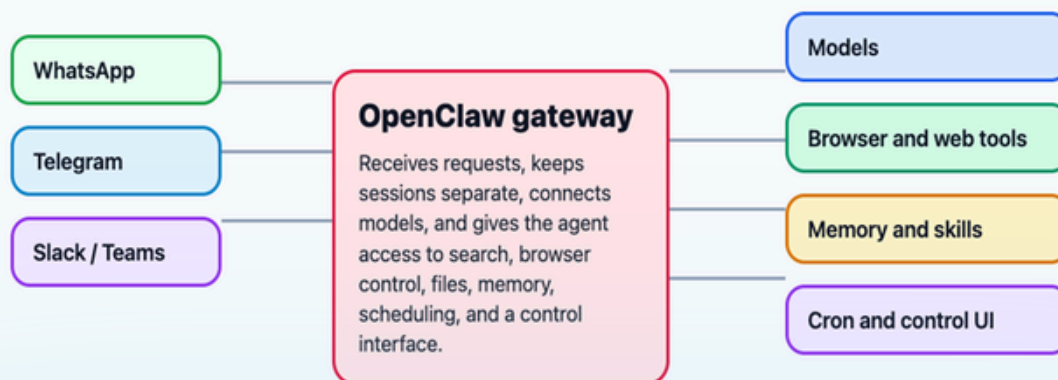


Figure 3. OpenClaw as a gateway agent. Original figure created for this article. Concept source pages: OpenClaw – Personal AI Assistant (<https://openclaw.ai/>).

Other commercial offerings for AI agents are more ready-to-use. Manus is designed for longer, delegated tasks within a controlled virtual environment. Perplexity is a well-recognised tool for in-depth research [6–8]. They are tools available in Hong Kong through franchised offerings or institutional subscriptions [9,10]. Numerous other options are available, each with its own strengths and design philosophy. For users without a technical background, these commercially developed products are generally more accessible and ready for immediate use.



June 2026

NEWSLETTER |

Hong Kong College
of Emergency
Medicine

Digital EM subcommittee: Agentic AI and OpenClaw

The cautious horizon

AI agents are particularly well suited to information-heavy and repetitive clerical work, representing a form of ambient automation that reduces the need for constant prompting. These systems can help professionals offload some of the cognitive load and distractions.

However, no agentic tool eliminates the need of human oversight. Under-polished AI tools can produce inaccurate outputs or omit critical details. The agent may respond erroneously in slightly different contexts. These shortcomings had limited the utility of agentic tools. The World Health Organization cautioned that AI systems, while promising across many areas of healthcare often “hallucinate” and contribute to automation bias, and introduce cybersecurity vulnerabilities if not properly governed [11]. Ongoing developments, such as chain-of-thought reasoning, retrieval-augmented generation (RAG), and knowledge-integrated architectures, alongside increasing model scale and computational power, aim to mitigate these issues but do not eliminate them entirely. Our cautious horizon is utterly important especially when AI agents begin to assist in patient care.

AI agents are only valuable if they are also safe. Unfortunately, the adoption of AI has increased attack surfaces and novel threats such as prompt injection. Malicious instructions embedded within seemingly benign documents can mislead AI systems into producing harmful outputs. Misunderstanding of user’s instruction would also lead to catastrophic results. These scenarios are especially dangerous when agents had become capable to interact with external tools. In the world of software development, incidents of AI coding agents inadvertently deleting entire codebases or databases was reported [12,13].

A cautious and staged approach is therefore essential, particularly in high-stakes environments. Effective safeguards include robust permission management, clearly defined privacy boundaries, sandboxed execution environments, comprehensive logging, human-in-the-loop approval mechanisms, and auditable activity trails. Most current AI agent frameworks already incorporate forms of supervised operation, where explicit human authorisation is required before executing higher-risk action beyond predefined boundaries. For now, the safest use for AI agents in clinical context is ambient AI tools and decision-making support helpers, rather than as fully autonomous agents.

The doctor’s second shift and more

With the current offerings and the limitations discussed above, the most realistic use cases for AI agents for doctors are primarily administrative and clerical work that constitute the “second shift”—the unseen administrative burden that continues after clinical hours. Some more mature AI capabilities, albeit non-agentic, are already in wider clinical use including real-time translation in consultations, AI-assisted clinical documentation (scribing), and automated generation of medical reports and referral letters. In public hospitals in Hong Kong, some of these functionalities are already accessible through the Hospital Authority’s HA Chat application or web interface.



Evidence already suggests meaningful efficiency gains from current natural language processing tools. Ambient AI scribes, for instance, have been reported to reduce clinical documentation time in emergency department settings by approximately 16–39% in real-world and simulated studies, and are associated with reductions in physician burnout [14,15]. One study further reported a 28% reduction in on-shift documentation time and a 16% reduction in total time spent on electronic health records [15]. Yet, a scoping review has highlighted issues such as transcription errors requiring manual correction, excessive note length, formatting inconsistencies, integration hurdles, and performance variability that has hindered broader adoption [16]. In Hong Kong, widespread bilingualism presents an additional challenge, particularly given the complexity of Cantonese for current language models.

In the near term, the most realistic integration of AI agents in healthcare limits to non-clinical workflow support, such as documentation and feature flagging. It can compile clinical notes, laboratory results, and medication changes to generate draft discharge summaries, while also flagging missing follow-up tasks or appointments prior to discharge. AI-enabled drug interaction and cross-sensitivity system is already adopted in public hospitals. On the research and audit side, more AI-enabled tools are being incorporated into clinical databases now to facilitate machine learning and advanced data-analysis. For more clinical applications, decision support system is increasingly popular. By integrating with clinical equipment, such as physiological monitors and point-of-care testing devices, and imaging or ECG interpretation systems, combining with validated clinical algorithms, clinicians can make more informed judgements. For in-patient settings, “clinical copilots” that can automatically summarize ward rounds, extrapolate laboratory trends, and estimate readiness for discharge are emerging. In all these examples, the AI systems perform supervised automation, where physicians remain central to final decision-making. By contrast, unsupervised systems that independently diagnose, prescribe, or determine patient disposition remain far from clinical reality. Beyond technical maturity, such applications raise unresolved questions around professional autonomy, medicolegal liability, and regulatory approval, all of which must evolve alongside the technology itself.

Figure 2. A supervised clinical agent workflow

The near-term value is mostly administrative. The doctor still makes the final decision.



Figure 2. A supervised clinical agent workflow. Original figure created for this article. Concept source pages: WHO guidance on large multi-modal models in health (<https://www.who.int/publications/i/item/9789240084759>) and JMIR emergency department ambient scribe study (<https://formative.jmir.org/2026/1/e80401>).



June 2026

NEWSLETTER I

Hong Kong College
of Emergency
Medicine

Digital EM subcommittee: Agentic AI and OpenClaw

Reflection

This passage is entirely written by the author, not even proofread by AI.

The author also has no affiliation with, nor made any investment into any of the mentioned companies or project developers.

While preparing for this article, a sudden and intrusive idea came up in my mind: what if I used AI agents to write about AI agents? Instead of running straight to GoogleLM for quick results immediately, I decided that I should try building and running my own agent to test its ability. While commercial products are obvious shortcuts, are similar tools not behind a paywall going to give us an equally seamless experience? Like how an AI agent would, I devised a plan, researched on OpenClaw and then began tinkering on my laptop. “Not without friction” is an understatement, as expected. The setup screen began with a disclaimer about it being a hobby project and advise against using it in production environments. While using AI agents felt like hopping on an Uber, working with the claw feels like driving a hobby project car. It runs, but it breaks down more often. The greatest limiting factor is probably me, as a non-technical person. We are also seeing people running an entire one-man business with them. Eventually I read about a contender -- The Hermes Agent, and I immediately tried. It is still not painless to fiddle with, the learning curve still existed but manageable. As the name of the Greek god of speed and communication suggests, it felt faster and was very responsive to my random questions on WhatsApp. I asked Hermes about knowledge in by Obsidian Notes vault, and I prompted it to give me heads up about the weather before each shift. It mostly responded with what I needed.

Eventually I decided to ask Hermes to write this article. The prompt and the raw output of can be found at <https://github.com/kellogcheung/agenticai-and-openclaw/>. The headings, the graphics, formatting and the list of citations and the file itself was the final output from the agent. I did go back and forth with a few times to refine the tone and format. If you bothered to read that article, it looked nothing like this one. It has a robotic writing style, a very jumpy logic and awkward headings. I almost rewrote it entirely. In the end, the whole process of setting up the agent to write and refine this article probably took me longer than writing it myself. It was certainly not a good advertisement for agentic AI, but it was a fun experiment for me.

The postmortem is simple: the computational power was limited, and the tools installed onto the agent were less polished. For this article, I ran the agent on GPT-5.4 via GitHub Co-Pilot API. It is not a small model but not the current best. Instead of the more sophisticated research skills written by Anthropic, I used the built-in research skill bundled with Hermes Agent, which is still in early beta. On the other hand, without optimised prompting techniques and well-engineered context, an agent is only as dumb as any unsupervised chatbot, hence the result. As the agent was advertised as self-improving, I hope that it would gradually consolidate its memory about my preferences as we continue to interact. Eventually it may really learn to work and communicate on my behalf, not just as an assistant, but a digital presence and proxy of the physical me.



June 2026

NEWSLETTER |

Hong Kong College
of Emergency
Medicine

Digital EM subcommittee: Agentic AI and OpenClaw

References

- [1] OpenAI. A practical guide to building agents. 2025. Available from: <https://openai.com/business/guides-and-resources/a-practical-guide-to-building-ai-agents/> Accessed 22 Apr 2026.
- [2] Anthropic. Building Effective AI Agents. 2024 Dec 19. Available from: <https://www.anthropic.com/research/building-effective-agents> Accessed 22 Apr 2026.
- [3] OpenClaw. OpenClaw – Personal AI Assistant. Available from: <https://openclaw.ai/> Accessed 22 Apr 2026.
- [4] Steinberger P. Introducing OpenClaw. OpenClaw Blog. 2026 Jan 29. Available from: <https://openclaw.ai/blog/introducing-openclaw> Accessed 22 Apr 2026.
- [5] Nous Research. Hermes Agent documentation. Available from: <https://hermes-agent.nousresearch.com/docs/> Accessed 22 Apr 2026.
- [6] Manus. Welcome - Manus Documentation. Available from: <https://manus.im/docs/introduction/welcome> Accessed 22 Apr 2026.
- [7] Perplexity. Introducing Perplexity Deep Research. 2025 Feb 14. Available from: <https://www.perplexity.ai/hub/blog/introducing-perplexity-deep-research> Accessed 22 Apr 2026.
- [8] Perplexity. Introducing Perplexity Labs. 2025 May 29. Available from: <https://www.perplexity.ai/hub/blog/introducing-perplexity-labs> Accessed 22 Apr 2026.
- [9] Hong Kong Metropolitan University. ZoomIt tool page mentioning Manus. Available from: <https://www.hkmu.edu.hk/oetools/zoomit> Accessed 22 Apr 2026.
- [10] HKT/csl. HKT x Perplexity Pro promotion. 2025. Available from: <https://www.hkcsf.com/en/hkt-perplexitypro-launch-2025/> Accessed 22 Apr 2026.
- [11] World Health Organization. Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models. 2025 Mar 25. Available from: <https://www.who.int/publications/i/item/9789240084759> Accessed 22 Apr 2026.
- [12] Beatrice Nolan. An AI-powered coding tool wiped out a software company's database, then apologized for a 'catastrophic failure on my part'. Fortune. Available from: <https://fortune.com/2025/07/23/ai-coding-tool-replit-wiped-database-called-it-a-catastrophic-failure/> Accessed 22 Apr 2026.
- [13] Mark Tyson. Claude-powered AI coding agent deletes entire company database in 9 seconds – backups zapped, after Cursor tool powered by Anthropic's Claude goes rogue. tom's Hardware. Available from: <https://www.tomshardware.com/tech-industry/artificial-intelligence/claude-powered-ai-coding-agent-deletes-entire-company-database-in-9-seconds-backups-zapped-after-cursor-tool-powered-by-anthropics-claude-goes-rogue> Accessed 27 Apr 2026.
- [14] Lipinski M, Smith E, Vaillancourt C. Human vs. artificial intelligence in medical charting: a comparative study in the simulated emergency medicine context. CJEM. Available from: <https://pubmed.ncbi.nlm.nih.gov/41917375/>
- [15] Preiksaitis C, Alvarez A, Winkel M, et al. Ambient Artificial Intelligence Scribe Adoption and Documentation Time in the Emergency Department. Ann Emerg Med. 2026;87(5):569-574. Available from: <https://pubmed.ncbi.nlm.nih.gov/41665590/>
- [16] Atiku S, Olakotan O, Owolanke K. Usability-Related Barriers and Facilitators Influencing the Adoption and Use of AI Scribes in Healthcare: A Scoping Review. J Eval Clin Pract. 2026;32(1):e70365. Available from: <https://pubmed.ncbi.nlm.nih.gov/41533034/>

June 2026



NEWSLETTER

Hong Kong College
of Emergency
Medicine

Upcoming Events

JCM JUN 2026

Date: 3 June 2026 (Wed)

Time: 5:30- 7:30pm

Venue: Seminar Room, G/F, Block A, Queen Elizabeth Hospital

Program

5:30-6:15pm Case Presentation by AED, TKOH

6:15-7:00pm "Revisiting Heart Failure - Basics and Beyond"

Dr MOK Ho Fung, Corbin, Resident, Department of Medicine & Therapeutics, PWH

7:05-7:30pm Discussion on OSCE by AED NDH



JCM JUL 2026

Date: 8 Jul 2026 (Wed)

Time: 5:30- 7:30pm

Venue: H702, 7/F Block H, Princess Margaret Hospital

Program

5:30-6:15pm Case Presentation by AED, TSWH

6:15-7:00pm Fragility Fracture of the Pelvis

Dr. CHAN Wai Ming Consultant

Department of Orthopedics and Traumatology, TMH

7:05-7:30pm Discussion on OSCE by AED PYNEH



Private Healthcare Facilities Ordinance (Cap. 633)

**Apply for a Clinic Licence or
Request a Letter of Exemption
for a Small Practice Clinic now**



智方便
iAM Smart



e-Licensing system



Join our briefing



Simple guide to
Cap. 633

Enquiry

3107 8451 (medical)
2631 1782 (dental)



衛生署

Department of Health

June 2026

NEWSLETTER |

Hong Kong College
of Emergency
Medicine

Upcoming Events



溪降活動日

ORGANIZED BY
HKSEMS & HKCEM

日期：28/9/2026 | 時間：0900-1700

地點：屏南石澗 | 名額：10人

費用：\$300(包保險及裝備)

查詢：Joanne

Tel: 28718875

Email: joanne@hkcem.org.hk

額滿即止

