BMJ Clinical Intelligence

Why the BMJ Knowledge Graph

Dr. Blackford Middleton

Dr. Kieran Walsh

Dr. Blackford Middleton is a Physician and Health Informatics Innovator committed to global transformation … curing what ails healthcare, improving healthcare one decision at a time. His special interests are in augmented clinical intelligence, knowledge representation and sharing, and quality measurement – all mediated through clinical information management systems.

He has deep experience in team building, innovation, and leadership roles in academe and industry. He was a professor of biomedical informatics, and/or of medicine, at Stanford, Harvard, and Vanderbilt Universities, and he held executive leadership roles at Stanford University Medical Center (CMIO), MedicaLogic/ Medscape (CMO), Partners Healthcare System (Corporate Director, Clinical Informatics R&D; now Mass General Brigham), Vanderbilt University Medical Center (CIO/Assistant Vice Chancellor), and at Apervita (Chief Informatics & Innovation Officer). At the BMJ, he serves as consultant for Digital Knowledge Products.

Dr. Kieran Walsh is Clinical Director at BMJ. He is the clinical quality and safety lead of the medical education and clinical decision support resources at BMJ.

He has a vast amount of experience in clinical decision support, online medical education, face-to-face delivery of medical education, and both summative and formative assessment. He has experience of using all of these in programs to strengthen health systems.

He is Adjunct Associate Professor in Teaching and Research at Monash University. He has published over 200 papers in the biomedical literature and has written four books on medical education.

Summary

Modern healthcare delivery faces multiple challenges, including delayed translation of new knowledge into clinical practice, unwarranted variability in healthcare delivery at the patient and population level, and lack of optimization of cost, quality, and value. Traditional clinical decision support has attempted to tackle some of these problems. However, to date it has largely been developed on rule-based systems. These have been helpful, but the rule base becomes difficult to manage at scale – for example when you have to manage hundreds of rules. And typically rules take a single disease perspective which will not work for patients with multimorbidity.

An alternative is a knowledge graph approach. BMJ Clinical Intelligence is our knowledge graph, and it has been designed to overcome the limitations of traditional rule-based systems. BMJ Clinical Intelligence is an evidence-based and continuously updated resource covering a range of medical specialties. Its structured representation of knowledge allows for a comprehensive understanding of relationships between medical entities, fostering better clinical decision-making.

The BMJ Clinical Intelligence approach involves knowledge engineering at scale, ensuring the translation of clinical guidelines into computable evidence.

With a team of clinical experts and informatics specialists, BMJ provides a solution that spans clinical medicine and population health. The knowledge graph is updated daily, offering a dynamic, comprehensive, and evidence-based tool for healthcare professionals.

BMJ Clinical Intelligence can be utilized in clinical decision support at the point of care and in population health analysis. Its knowledge graph can identify deviations in patient care, suggest interventions, and facilitate early diagnosis. In population health, the graph enables comprehensive monitoring, analysis, and early intervention, ensuring efficient resource utilization.

The knowledge graph's advantages over traditional guidelines lie in its ability to represent an entire domain, providing a holistic view of patient care. It complements large language models, enhancing clinical reasoning and decision support. The continuous updating and scalability of knowledge graphs offer a dynamic and comprehensive approach to healthcare knowledge management and clinical care. BMJ Clinical Intelligence provides knowledge that is relevant to every decision in healthcare, and that will improve each decision – from the point of care to population health.

Problems in healthcare

Modern healthcare delivery faces multiple challenges, including delayed translation of new knowledge into clinical practice, unwarranted variability in healthcare delivery at the patient and population level, lack of optimization of value, cost, and quality, a need for improved patient and provider experience, and the need for payment reform.

Delayed translation of new knowledge to clinical practice is caused in part by the ongoing knowledge explosion in medicine. This continues apace and it is extraordinarily difficult for general internists to keep up with everything in internal medicine or primary care. Sub-specialists have a narrower but deeper domain and so face their own challenges. Suffice to say, the rate of discovery, publication, and learning results in a delayed translation of new evidence into practice for all clinicians. Some have estimated that it takes up to 17 years for a new notable and clinically valid innovation to be generally adopted and applied in healthcare (Balas and Boren 2000).

Unwarranted variability in healthcare delivery is another thorny problem. Clinicians today will always try to practice at the top of their license all the time. But it is really difficult to have exactly the right knowledge available for the right decision in the right place and at the right time in all settings. And when that current knowledge is not available, the result is unwarranted variability where the patient may not experience the care that they should. There may also be a lack of personalized care that is tailored to each unique patient.

Research has shown that evidence-based medicine is only applied in clinical practice about 54% of the time

(McGlynn EA et al).

Only when we can leverage the best evidence in the right context will we be able to reduce unwarranted variation in care.

There is also a lack of optimization of value, cost, and quality and a need for improved patient and provider experience – the Quadruple Aim. There are multiple causes of these problems but one fundamental problem underlying many of them is that the clinician does not always have access to the best evidence for each and every decision that they need to make. And, sometimes they do not have the ability to apply that evidence at the point of care in the patient's particular context and setting.

Lastly, there is the issue of payment reform. Many payers are interested in reforming the ways in which providers are paid for services. They want to incentivize providers to do more of the right things that they should be doing, and less of the wrong things that they should be avoiding. This requires knowing first of all, what is the right thing and what is the wrong thing to do clinically. Having that knowledge available and applied at the point of care in real-time should optimize physician behavior and decision-making in a way that will benefit payers, providers, and patients.

Many of the challenges relate to optimizing the clinical encounter in all the settings where it might occur and also overcoming current problems with clinical decision support.

Problems with clinical decision support systems

Clinical decision support systems have been around for some time. They have been developed in research and commercial settings, and are largely delivered by commercial entities. But there are challenges associated with clinical decision support systems. Just a few of these challenges are outlined below.

> Most clinical decision support has been developed as rule-based systems. These have been helpful in the past, but the rule base may become difficult to manage at scale – for example when you have to manage hundreds of rules.

There can be rule conflict and resultant confusion when multiple rules are operating that do not agree with each other. And rules necessarily take a single disease perspective which will not work for patients with multimorbidity. This is a big problem when up to one-third of the adult population has multimorbidity. Put simply, the clinical workforce is faced with an aging population with a growing number of complex interrelated problems.

Other basic problems include the need for too much typing and alerts that interrupt the clinical workflow. Such alerts are often out of context and lead to alert fatigue and user dissatisfaction. In some cases, users switch off alerts completely as they may not be viewed as helpful, are not actionable, or may adversely impact their clinical workflow.

Lastly, there is the issue of the management and maintenance of large knowledge assets that are part of clinical decision support systems – given the ongoing evolution of clinical knowledge and changing controlled medical terminology standards. These are simply too much for any single institution or provider to take on. Even large academic and clinical institutions struggle to build knowledge management teams to manage all of these clinical knowledge assets that are being deployed across multiple sites.

These problems are not new – they have been known about for many years. Clinical decision support systems have struggled to overcome them. Some efforts have resulted in even more rules and alerts – the very problem we are trying to solve.

What is the BMJ Knowledge Graph?

A knowledge graph is a structured representation of information that captures the relationships between entities. It is a way to organize knowledge and data, and it highlights connections among various pieces of information.

BMJ Clinical Intelligence is the knowledge graph designed by BMJ. It is a set of continually updated, evidence-based, and structured knowledge assets. Our knowledge assets are created from evidence and guidelines, covering both single conditions and complex comorbidities.

Updated daily using robust methodologies, our clinical information is encoded and converted into computable form in a knowledge graph which drives algorithms that augment clinical decision-making.

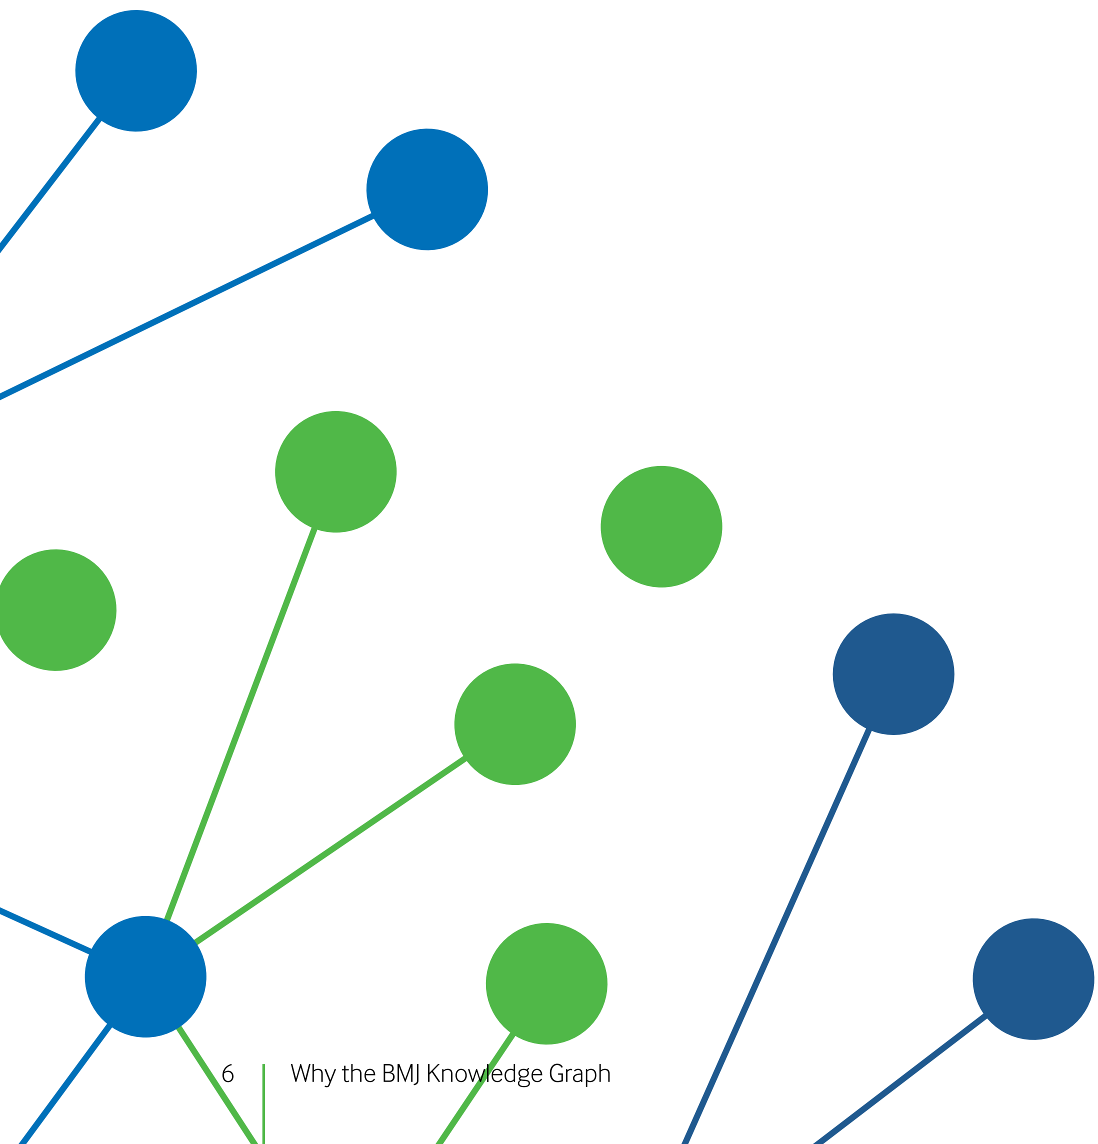

Why the BMJ Clinical Intelligence knowledge graph approach?

The problems with healthcare and clinical decision support outlined above mean that we need a fundamental rethink of how we deliver knowledge to healthcare professionals.

First and foremost is the need for knowledge engineering at scale. This will enable the translation of clinical guidelines or best practice documents and protocols into computable evidence that can be used at the point of care. What is needed to achieve this is a multidisciplinary team and processes that require both clinician expertise and consensus to arrive at clinical guidelines, and clinical informatics expertise to translate and encode guidelines into computable evidence.

BMJ is well-positioned to lead in this work. BMJ has published in the field of medicine and healthcare since 1840. It is a world-class medical publisher with decades of experience in publishing top-tier clinical journals that are known worldwide. BMJ has also been providing digital clinical decision support for over fifteen years. It has now built an expert knowledge engineering team with deep experience in the translation process from guidelines to computable evidence at scale across whole domains – from clinical medicine to population health. The result is computable evidence that can be applied at the point of care and in population health management for each and every clinical decision made.

The BMJ Clinical Intelligence knowledge graph is built upon knowledge assets that are evidence-based and have a clear provenance that cites the source materials. There is a continuous updating process to ensure that the knowledge is current. BMJ has built a factory that keeps all the knowledge assets updated and at the same time in synchrony with the translated computable evidence.

The knowledge assets are also comprehensive: they cover over thirty medical specialties across the breadth of primary, secondary, and tertiary care. They cover each topic in-depth and cover the entire clinical workflow at the point of care – including screening, triage, diagnosis, differential diagnosis, investigations, treatment, follow-up, and healthcare maintenance. The knowledge assets work at the population health level as well – they include comprehensive population health monitoring, intelligent analysis, and early intervention that can enable better screening and diagnosis, prompt and comprehensive treatment, and primary and secondary prevention.

Individual institutions would struggle to do this at scale. BMJ has built the capability to do this and to collaborate with partners who can enable implementation at the point of care and in population health management.

How might BMJ Clinical Intelligence be used?

BMJ Clinical Intelligence can be used in multiple contexts. It can be used in clinical decision support at the point of care via web services or as a SMART on FHIR application for individual patients. It can also be used as a knowledge resource in population health to address the care of cohorts of patients.

> At the point of care, the graph can be queried to look for deviations from a patient's expected course – for a care gap or untoward trend. And in turn, once it detects a deviation, it can suggest an early intervention to rectify the patient's care.

It can help in cases of diagnostic error or diagnostic delay. Patients might have a delayed diagnosis which can be impactful on their clinical course. The knowledge graph can be used to help accelerate the differential diagnosis reasoning process to augment the clinical reasoning of the physician and so avoid error or delay.

BMJ Clinical Intelligence can be used in population health analyses as well. For example, it could be used on a whole population and then a cohort of that population who have a given condition such as diabetes and finally a section of that cohort with diabetes who are not proceeding as expected in their care journey – they might exhibit care gaps or untoward trends. For example, the knowledge graph might provide the clinical context and prioritization of patients identified as frequent flyers or high utilizers of emergency care services. Healthcare providers may then call in such patients preemptively to keep them on track and to prevent them from becoming acutely unwell. Thus, limited care resources are directed to where they are most needed – resulting in accurate patient risk stratification and more cost-effective resource allocation.

Today knowledge representation at scale is best done as a knowledge graph. This is true of other industries, not just healthcare. Knowledge graphs are supporting drug discovery in the pharmaceutical industry, fraud detection in the banking industry, and even space exploration at NASA. In this case, our knowledge graph is a set of diseases and findings that are all interrelated to each other to give a complete picture of a domain of practice. Knowledge graphs allow us to do decision support in novel ways. They allow us not only to do individual condition-specific decision support for a particular disease, they also allow us to look at a disease and find all of its neighbors and comorbidities to which the disease might be connected, and to find out how all of those relate to the patient's overall state.

In this way, we can augment clinical reasoning in a fashion that is not possible with rule-based approaches for single conditions.

The advantages of BMJ Clinical Intelligence compared to rule-based approaches

In clinical medicine, we have had decades of experience with rule-based approaches. Many institutions and vendors have delivered those kinds of resources and tools. But they are difficult to manage at scale. Organizations need dozens of people to be involved in maintaining large bodies of rules and when you get into hundreds of rules, it becomes difficult to test and determine if they are all working appropriately in synchrony. The knowledge graph approach as opposed to the simple rule-based approach allows us to create a single model for a whole domain of practice be it in clinical medicine or population health.

Knowledge graphs allow us to reason in different ways as compared to simple rule-based approaches. A knowledge graph allows us to use semantic reasoning and look across the graph for relationships between diseases and other diseases, between diseases and findings, between findings and other findings, between findings and drugs, and even between drugs and diseases.

> This knowledge graph approach now undergirds most of the modern big tech stacks around the world. It is behind some of the best mapping systems, social media, and search and retrieval systems.

This approach is massively scalable and it is exciting to consider how knowledge graphs may work with, leverage, or complement large language models to define relationships, extract features, and make associations. These will all result in deeper insights that will benefit both patients and physicians.

Knowledge graphs also give us the ability to project the entire patient data set against the graph and to get a better picture of the patient – a more comprehensive picture than what might emerge from using forward and backward chaining through rule bases. When we see the whole patient state reflected against the graph, it can tell us for example that we expected to know a certain fact about the patient but we do not see it or we do not see that a relevant procedure or investigation has been considered – as it is not seen in the data. This is a patient-centered, data-driven approach – projecting the data against the knowledge graph and getting a snapshot of the patient's current state, and telling the end user who may or may not be familiar with that patient what are reasonable expectations for their care delivery or care journey.

Another feature of knowledge graphs is their ability to handle much more complexity than simple rule-based approaches. Rules tend to be relatively condition-specific and so it is difficult to express the complexity of multimorbidity with rule-based systems. So knowledge graphs allow us to see a more comprehensive picture not only of the care domain but of the patient state that is reflected in the care domain.

Rules can lead to a lot of redundant alerts or even inappropriate alerts. It is hard to tailor the alert appropriately to a specific user. For example, a cardiologist may wish to have a different set of alerts than a general practitioner. And a pediatric cardiologist may wish to have different alerts than a general cardiologist – so it quickly becomes complicated and difficult to scale. An alternative is to include models of the end users themselves in the knowledge graph and fine-tune decision support that is relevant to them – in essence both patient and provider-centric decision support.

Updating can also be a challenge in rules-based systems. Every time the evidence changes, it needs to be analyzed and potentially incorporated into a range of rules. And because rules are specific, every time one part of the evidence changes, the rules need to be reviewed to see if they need to change as well. The same may occur with changes to codes for controlled medical terminologies. However, knowledge graphs can be updated continuously – and BMJ has a team of expert clinical academics and consultants from around the world (the majority are US-based) to survey the literature and guidelines and to apprise the BMJ when new practice-changing evidence needs to be incorporated into its clinical information and then translated into computable evidence by the clinical informatics team.

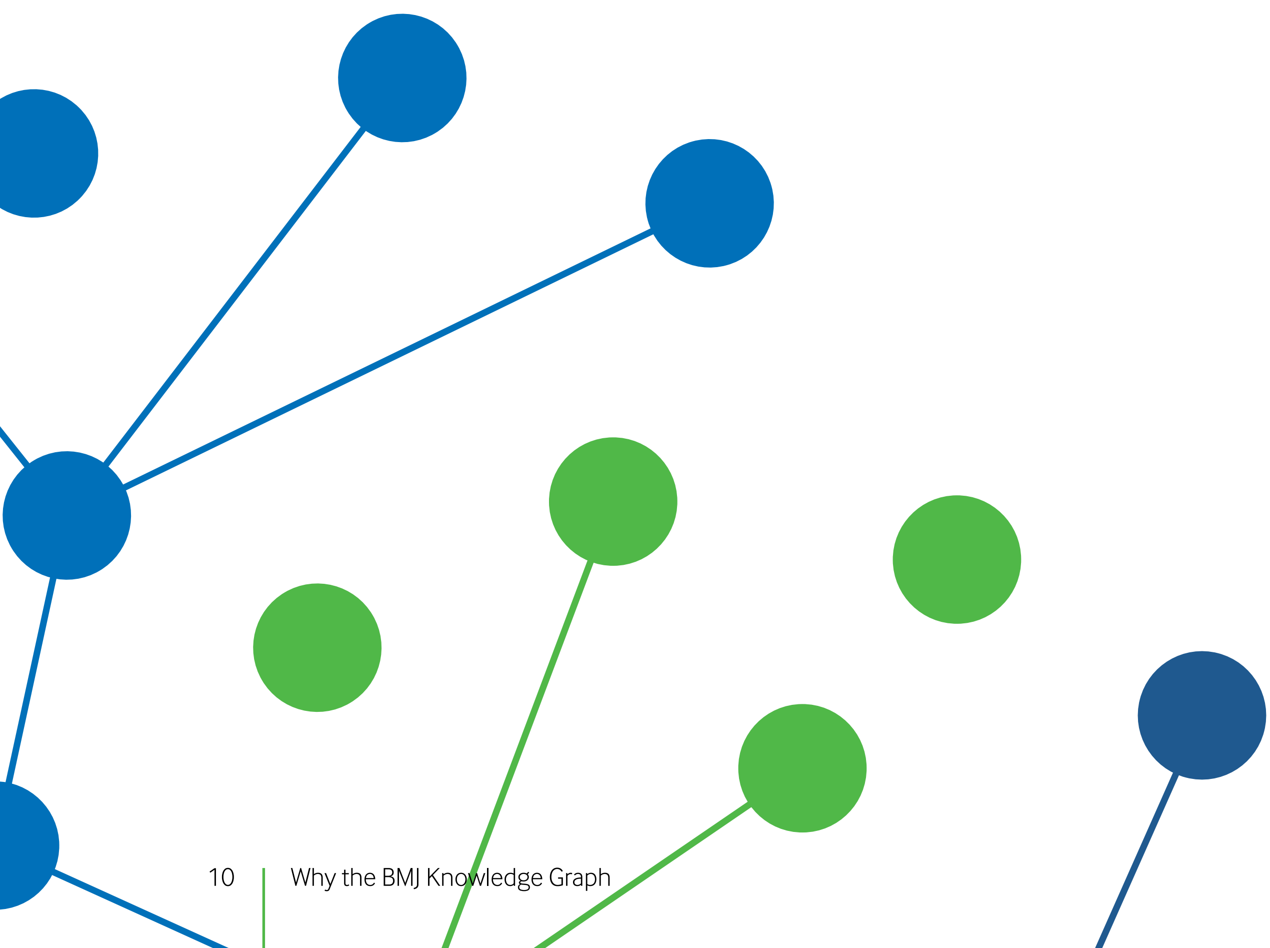

How BMJ Clinical Intelligence can work with and complement large language models

Knowledge graphs may leverage, work with, or complement large language models both to extract features, define relationships, and make associations, as well as to potentially fuse knowledge graph approaches and large language models to improve prompt engineering.

In large language models, getting the prompt or question right is all important. If you do not get the question right, you are never going to get the answer that you need. So knowledge graphs can help frame the question or prompt by providing continually updated domain specific context and information that can then be submitted to a large language model. The submitted question should be much more likely to get a context-specific, relevant and appropriate response that will answer the question directly and thus augment clinical reasoning appropriately.

Knowledge graphs might also actually help to discover new insights that previously were not seen or known about – so discoveries are an exciting part of this. The result should ultimately be a more holistic understanding or complete picture of the patient state within a domain of knowledge. This in turn should make it easier to augment clinical reasoning appropriately and help healthcare professionals make better decisions.

The advent of large language models operating at scale with trillions of nodes or parameters that can be used to infer what might be the next word in a chain of thought means we can really start to emulate clinical reasoning in many interesting ways. But the challenge is the notion of hallucinations or fabrications that may result when using a large language model. Sometimes it appears that the large language model just attempts to fill in a gap and so it makes stuff up. This is of course completely unacceptable in clinical decision-making – we would not want a large language model to guess, imagine, or fabricate a drug or test for a particular condition.

New research is looking at how large language models and knowledge graphs can complement one another. And the research is showing that they can do so in a variety of interesting ways. For example, the learning from a large language model might be guided by structured knowledge coming from the knowledge graph to make the resulting large language model far more robust and allow access to newer or updated information held in the knowledge graph.

The knowledge graph can also assess the output of large language models to validity check or to look for hallucinations or fabrications. Using a knowledge graph with a large language model also allows attribution of the references and source material. And the knowledge graphs and large language models can work together in a fused way – so-called fusion models where parts of the inference might be done by a large language model and parts done by a knowledge graph, with both working in synchrony.

What are the advantages of our approach compared to traditional guidelines?

Traditional guidelines typically cover the diagnosis and management of a single disease. Knowledge graphs are more like a whole encyclopedia – a comprehensive domain model that is not single disease-specific or single drug-specific. Knowledge graphs allow us to model an entire domain. We can fuse knowledge graphs together from different domains and then they will be more comprehensive and more amenable to being continuously updated. They can allow us to see known and even discover unknown relationships between different combinations of concepts as diverse as diseases, clinical findings, and drugs. This comprehensive graph-based connected approach means that the reasoning capabilities of knowledge graphs will much more closely emulate the thinking of the physician.

Today we can no longer see the patient before us as a single disease entity. We have to think about the patient in their full context; not only the multiple diseases and comorbidities which may be at play but the social factors and social determinants of health and support structures that can impact the patient and the care process as well. In fact, even payment reform may influence what we do and how we do it and ideally incentivize more of the right things and fewer of the wrong things and at the same provide continuous feedback to the clinician.

The aim is that the clinician end-user has a delightful experience using a knowledge graph-based augmented clinical reasoning system and that they will be in a continuous learning mode as a result of receiving continuous feedback.

If we do this for each and every doctor, we will be truly enabling a learning health system. The learning health system in this context is essentially the idea that we learn from real-world evidence as well as clinical guidelines. We model knowledge and make it available in computable form. We then deliver that to the point of care and observe, monitor, and update to improve the knowledge assets as they are used – enabling a virtuous learning cycle.

User experience of a knowledge graph-based augmented clinical reasoning system is also important. Sometimes we aim too low in our goals for user experience in decision support – we plan for the software to be usable or acceptable. But with knowledge graph-based augmented clinical reasoning, we should be aiming to make the physician's experience delightful – so that it is simple and fun for them to do the right thing and to support continuous learning in practice. In the same way that we can project patient data against the knowledge graph, it is also possible to project clinician experience and training data. The linked data and shared ontology approach allow us to improve user experience in a more powerful way.

We can use the context of the clinician to tailor what is presented. For example, factors such as how many years the clinician is post qualification; what they have been shown recently; whether they already overridden this type of recommendation, their speciality; and whether this recommendation is outside their specialist knowledge could all influence what is most valuable to display given the user's limited opportunity to take in new information. Knowledge graph-based augmented clinical reasoning should mean nudge decisions where you are fitting into the moment when the clinician thinks – I might have an information need or a knowledge gap – a 'teachable moment'. And they ask themselves, do I even bother to try to close the knowledge gap or do I just move on? If it is already halfway closed for the clinician and all they have to do is acknowledge it, then that is a win-win – for clinicians and for their patients.

At a more fundamental level, even the largest guideline producers can publish a limited number of guidelines on a limited number of diseases whereas knowledge graphs can cover a far wider range of conditions and in considerably more depth. In addition, guidelines are usually updated at fixed intervals whereas knowledge graphs can be subject to continuous updating. Another problem is that evidence and subsequent guidelines don't exist for all areas of clinical practice. This is especially true for diagnosis-focused recommendations, and those considered expert opinion best practices. A knowledge graph can scale to include this best practice-based consensus as well as the latest evidence-based recommendations. Lastly, many guidelines are quite long so it can be difficult to quickly find the exact nugget of knowledge that you need. In contrast, the knowledge graph approach enables Osheroff's five rights of clinical decision support: the right information, to the right person, in the right format, through the right channel, at the right time in the workflow.

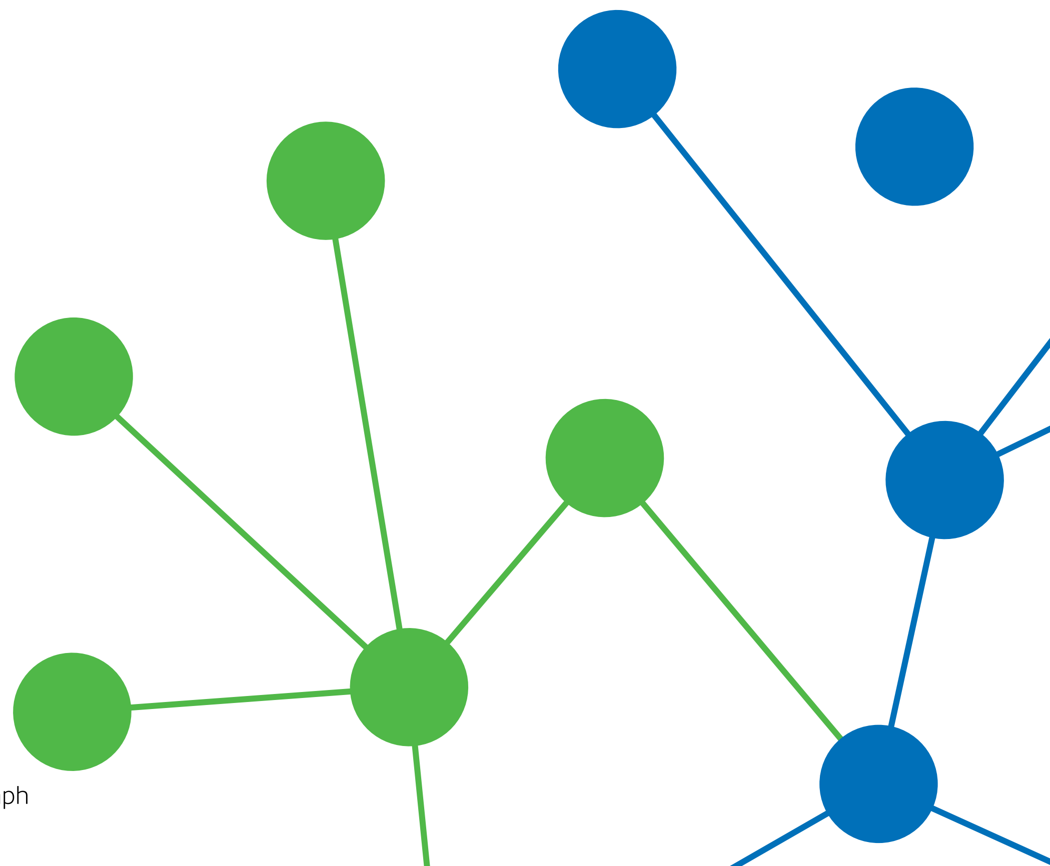

Best care everywhere

All patients want to receive the best possible care. And all healthcare professionals want to be able to deliver this, as indeed do all payers, providers, and platforms supporting care. To provide the best possible care, we need to provide the best knowledge at the time of clinical decision-making. We need to provide knowledge that is relevant to each and every decision in healthcare and that will improve each and every decision. We will then be able to transform healthcare – one decision at a time at first – but with a long-term vision of doing this at scale. It is a moral obligation to deliver the best possible care to all our patients.

It is now becoming a financial obligation to deliver best possible care and also best value for care with appropriate and optimal reimbursement. And of course, it is a clinical obligation to do what is right for each patient. One of the main gaps in achieving this is the gap between what is known in the world and what a clinician can access at the point of care. That knowledge gap is what we are fundamentally addressing with BMJ Clinical Intelligence so that each and every practitioner will have available to them the latest and greatest knowledge tool to use in care delivery. Ultimately when that gap is filled, we will all benefit.

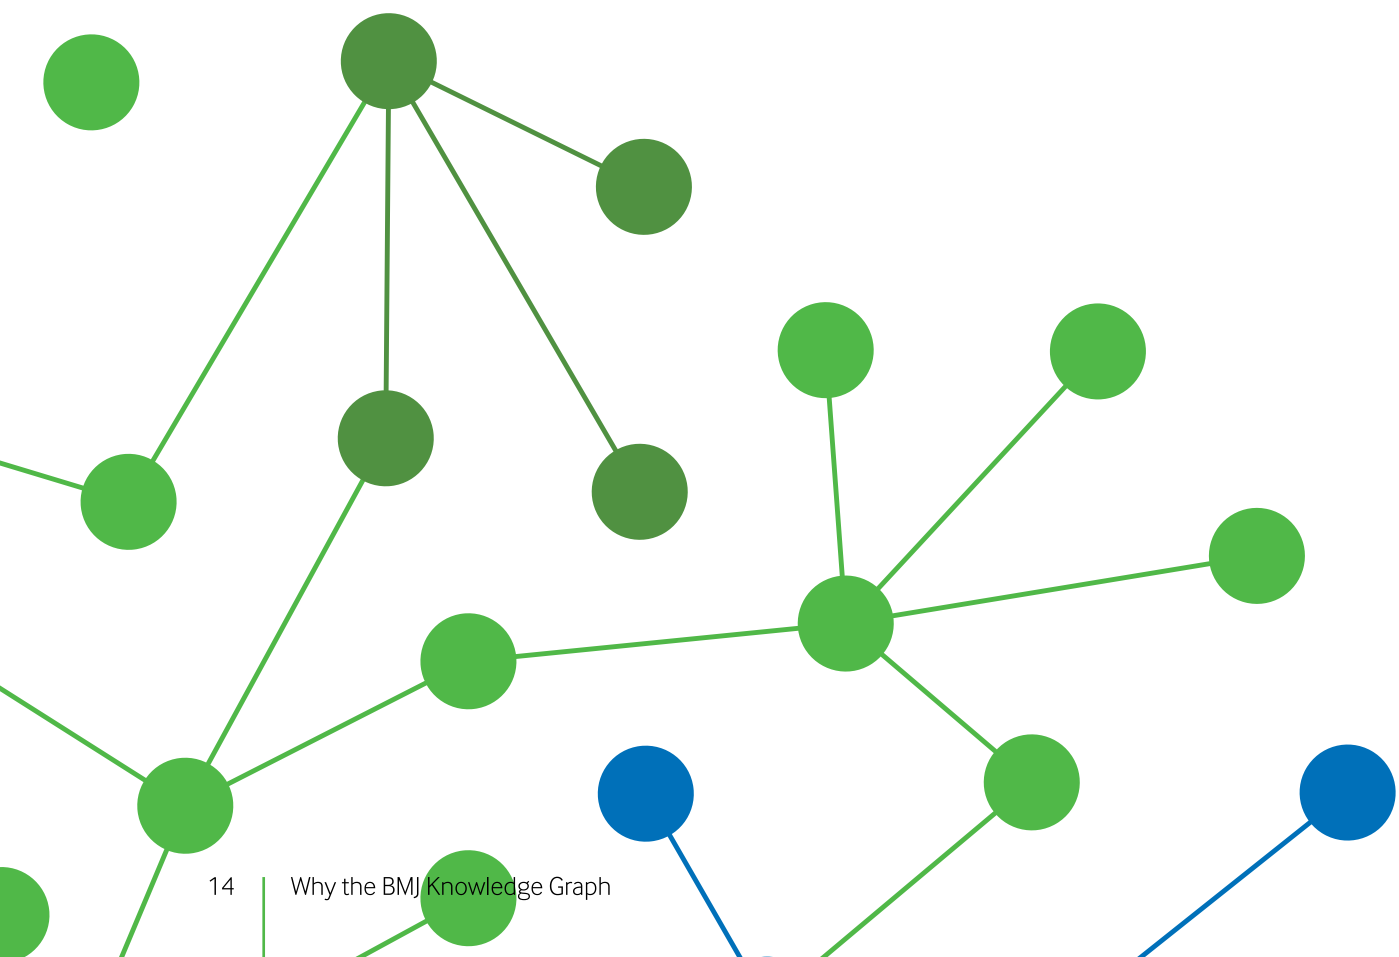

References / further reading

Balas, E. A. and S. A. Boren (2000). "Managing Clinical Knowledge for Health Care Improvement." Yearbook of Medical Informatics 09(01): 65-70.

BMJ Best Practice
https://bestpractice.bmj.com/info/

Bright TJ, Wong A, Dhurjati R, Bristow E, Bastian L, Coeytaux RR, Samsa G, Hasselblad V, Williams JW, Musty MD, Wing L, Kendrick AS, Sanders GD, Lobach D. Effect of clinical decision-support systems: a systematic review. Ann Intern Med. 2012 Jul 3;157(1):29-43.

Campbell JM, Umapathysivam K, Xue Y, Lockwood C. Evidence-Based Practice Point-of-Care Resources: A Quantitative Evaluation of Quality, Rigor, and Content. Worldviews Evid Based Nurs. 2015 Dec;12(6):313-27.

Doctors speak out about unnecessary care as cost put at $800bn a year
https://www.bmj.com/press-releases/2012/10/02/doctors-speak-out-about-unnecessary-care cost-put-800bn-year

Fox J, Khan O, Curtis H, Wright A, Pal C, Cockburn N, Cooper J, Chandan JS, Nirantharakumar K. Rapid translation of clinical guidelines into executable knowledge: A case study of COVID-19 and online demonstration. Learn Health Syst. 2020 Jul 14;5(1):e10236.

Improving Diagnosis in Health Care
https://nap.nationalacademies.org/catalog/21794/

Kendrick AS, Sanders GD, Lobach D. Effect of clinical decision-support systems: a systematic review. Ann Intern Med. 2012 Jul 3;157(1):29-43.

Kim S, Kim EH, Kim HS. Physician Knowledge Base: Clinical Decision Support Systems. Yonsei Med J. 2022 Jan;63(1):8-15.

McGlynn EA, Asch SM, Adams J, et al. The quality of health care to adults in the United States. N Engl J Med. 2003; 348: 2635-2645.

Medication Without Harm
https://www.who.int/initiatives/medication-without-harm

DE, McDonald KM, Meltzer DO. How much diagnostic safety can we afford, and how should we decide? A health economics perspective. BMJ Quality & Safety 2013;22:ii11-ii20.

Platt JE, Solomonides AE, Walker PD, Amara PS, Richardson JE, Middleton B. A survey of computable biomedical knowledge repositories. Mobilizing Computable Biomedical Knowledge Trust and Policy Work Group. Learn Health Syst. 2022 Jun 3;7(1):e10314.

Rayman G, Akpan A, Cowie M, Evans R, Patel M, Posporelis S, Walsh K. Managing patients with comorbidities: future models of care. Future Healthc J. 2022 Jul;9(2):101-105.

Richardson JE, Middleton B, Platt JE, Blumenfeld BH. Building and maintaining trust in clinical decision support: Recommendations from the Patient-Centered CDS Learning Network. Learn Health Syst. 2019 Dec 11;4(2):e10208.

Rubinstein PF, Middleton B, Goodman KW, Lehmann CU. Commercial Interests in Continuing Medical Education: Where Do Electronic Health Record Vendors Fit? Acad Med. 2020 Nov;95(11):1674-1678.

Sikka, R., et al. (2015). "The Quadruple Aim: care, health, cost and meaning in work." BMJ Qual Saf 24(10): 608-610.

Sordo M, Tokachichu P, Vitale CJ, Maviglia SM, Rocha RA. Modeling Contextual Knowledge for Clinical Decision Support.

AMIA Annu Symp Proc. 2018 Apr 16;2017:1617-1624. eCollection 2017.

Tao L, Zhang C, Zeng L, Zhu S, Li N, Li W, Zhang H, Zhao Y, Zhan S, Ji H. Accuracy and Effects of Clinical Decision Support Systems Integrated With BMJ Best Practice-Aided Diagnosis: Interrupted Time Series Study. JMIR Med Inform. 2020 Jan 20;8(1):e16912.

Vaucher C, Bovet E, Bengough T, Pidoux V, Grossen M, Panese F, Burnand B. Meeting physicians' needs: a bottom-up approach for improving the implementation of medical knowledge into practice. Health Res Policy Syst. 2016 Jul 18;14(1):49.

Walsh K, Wroe C. Mobilising computable biomedical knowledge: challenges for clinical decision support from a medical knowledge provider. BMJ Health Care Inform. 2020 Jul;27(2):e100121.

Wang H, Zu Q, Lu M, Chen R, Yang Z, Gao Y, Ding J. Application of Medical Knowledge Graphs in Cardiology and Cardiovascular Medicine: A Brief Literature Review. Adv Ther. 2022 Sep;39(9):4052-4060.

Williams M, Bray BE, Greenes RA, McCusker J, Middleton B, Perry G, Platt J, Richesson RL, Rubin JC, Wheeler T. Summary of fourth annual MCBK public meeting: Mobilizing computable biomedical knowledge-metadata and trust. Learn Health Syst. 2021 Dec 22;6(1):e10301.

Williams M, Richesson RL, Bray BE, Greenes RA, McIntosh LD, Middleton B, Perry G, Platt J, Shaffer C. Summary of third annual MCBK public meeting: Mobilizing computable biomedical knowledge-Accelerating the second knowledge revolution. Learn Health Syst. 2020 Dec 23;5(1):e10255.

Zhang D, Xiao L, Duan J, Chang X, Walsh K, Sandars J, Brown J, Dang X, Shen W, Du J, Cao Y. Understanding online self-directed learning using point of care information systems (POCIS): A plot study using a capability approach perspective. Med Teach. 2022 Dec;44(12):1413-1419.

Contact us

Contact us to learn how BMJ Clinical Intelligence can drive better clinical decisions at your organization.

W: clinicalintelligence.bmj.com
E: clinicalintelligence@bmj.com

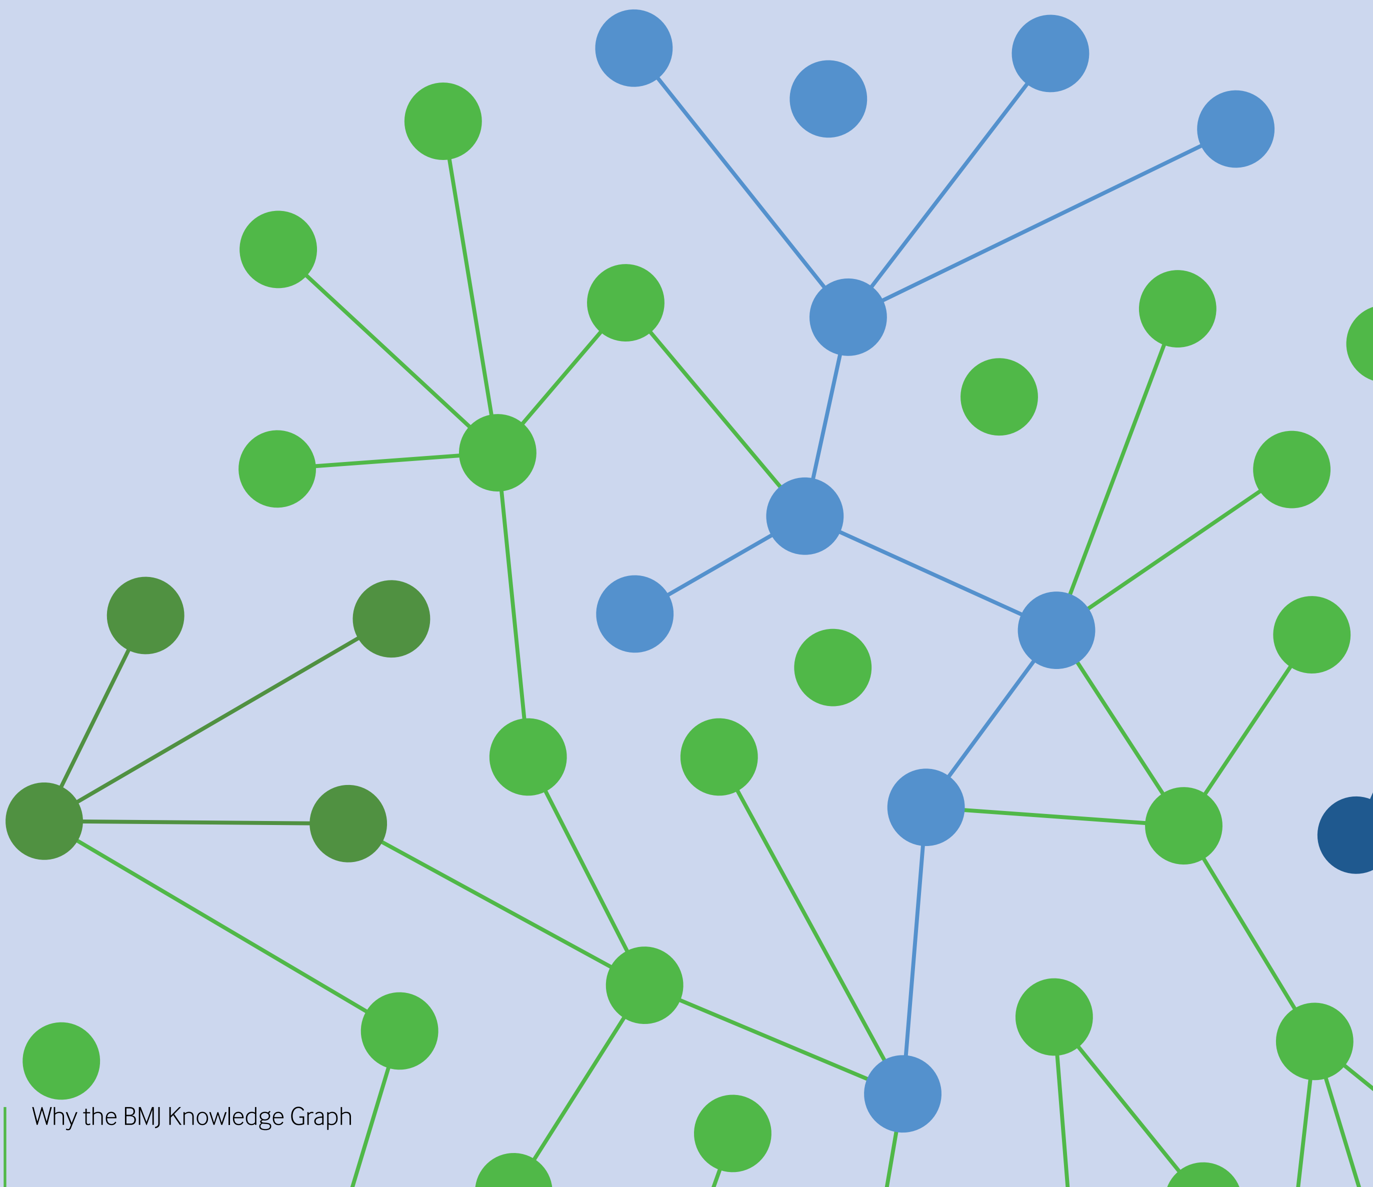

Notes

BMJ Clinical Intelligence

clinicalintelligence.bmj.com