

July 2026

Negotiating to Work in Partnership on AI:

The 2025 Agreement Between
Kaiser Permanente and
the Alliance of Health Care Unions

By Thomas A. Kochan, Erin L. Kelly, and Arrow Minster



About the Authors

Thomas A. Kochan is the George Maverick Bunker Professor of Management, Emeritus at the MIT Sloan School of Management and the Institute for Work and Employment Research (IWER). He is one of the authors of the book *Healing Together: The Labor-Management Partnership at Kaiser Permanente* (Cornell University Press, 2009).

Erin L. Kelly is the Sloan Distinguished Professor of Work and Organization Studies at the MIT Sloan School of Management, where she is Co-Director of the MIT Institute for Work and Employment Research (IWER).

Arrow Minster is an Assistant Professor of Management at the Lam Family College of Business at San Francisco State University.

Acknowledgments

Support for this research was provided by the Worker Empowerment Research Fund, the MIT Sloan Health Systems Initiative, Omidyar Network, and the MIT Institute for Work and Employment Research (IWER). We thank Martha Mangelsdorf for her excellent editing and Sasha Pedro for the creative design of the report. The views expressed are solely those of the authors.





TABLE OF CONTENTS

Executive Summary	4
Introduction	9
Background	10
How Participants Understood the Problem and Articulated Their Interests	19
<i>Patient and Mission-Centered Interests</i>	20
<i>Job Quality and Job Security Interests</i>	21
<i>Innovation Interests</i>	22
<i>Process-Centered Interests</i>	23
<i>Differing Understandings of “Inception”</i>	23
Developing Both Top-Down and Bottom-Up Strategies	25
Negotiations Over Economic Issues	29
Next Steps	31
Broader Implications of This Case	34
Moving Forward	38
Appendix	39
Notes	43

Executive Summary

Introduction

The 2025 negotiations between Kaiser Permanente (KP) and the Alliance of Health Care Unions (the Alliance) produced a landmark agreement on artificial intelligence and other related emerging technologies. This agreement represents the first comprehensive effort by a major healthcare employer and its union partners to integrate worker voice into all phases of AI decision-making—from problem identification through design, implementation, training, and transition assistance.

This report documents what was achieved, how it was accomplished, and what will be required to realize the agreement's full potential in practice.

Why This Matters

The fundamental challenge facing organizations implementing AI is not so much technological as it is organizational and human. As a 2025 report from the National Academies of Sciences, Engineering, and Medicine stated, the effect of AI on “worker and societal welfare necessarily depends on the legal, regulatory, and bargaining regimes in place.” Without the right structures, the report noted, the market may not produce outcomes that are beneficial to society.

KP and the Alliance recognized that for AI to deliver benefits to patients, employees, and the organization, workers needed a genuine voice in shaping how these powerful tools are developed and used. Their agreement provides a practical framework for making that commitment operational.

How the Parties Prepared

Both KP and the Alliance began preparing for negotiations in 2024, recognizing that AI would be a central issue. The parties took several intentional steps:

Building Understanding. In the fall of 2024, Thomas Kochan, one of the authors of this report, presented a framework for worker involvement in AI decision-making to Alliance leadership and then to the Joint KP Management and Alliance Executive Committee. These presentations introduced an integrated strategy for four aspects of technology implementation (problem identification, design, workforce education and training, and transition assistance).

Creating a Dedicated Subgroup. Rather than address AI in the context of traditional position-based bargaining, the parties agreed to form a subgroup of labor and management representatives using interest-based bargaining (IBB) methods. This approach focused on understanding underlying interests rather than stated positions, creating space for collaborative problem solving on an inherently complex issue.

Pre-Bargaining Engagement. In February 2025, approximately 300 Alliance leaders and 80 KP leaders convened for a joint pre-bargaining conference. This unusual event—rare in collective bargaining—allowed the parties to hear from each other about expectations for negotiations and engaged participants in workshops exploring their experiences with and concerns about AI. The conference revealed significant variation in views: some feared job loss, others already used AI to support their work, and many sought clarity about how they might shape AI's role in their jobs.

Understanding the Parties' Interests

Through more than 150 hours of subgroup deliberations, the parties articulated and explored four broad categories of interests:

Patient and Mission-Centered Interests. Both sides shared a commitment to high-quality patient care and organizational mission. Concerns about patient data stewardship, trust (particularly in communities with historical reasons for skepticism), and maintaining the “human touch” in healthcare were central to discussions.

Job Quality and Meaningful Work. Labor emphasized that AI should ease burdensome tasks and create opportunities for professional growth, not undermine autonomy or reduce workers to routine tasks. There was broad agreement that maintaining human decision-making authority and protecting professional judgment were essential, though the parties held somewhat different views about what this meant in practice.

Innovation and Organizational Agility. Management stressed the need to move quickly with AI adoption to remain competitive and serve patients effectively. Both sides recognized, however, that inadequate worker involvement often creates the very delays and implementation problems management sought to avoid.

Clarity, Transparency, and Partnership Processes. Both labor and management representatives valued clear structures, effective communication, and accountability. Labor representatives emphasized that trust and effective implementation require transparency about how AI is used, particularly in HR functions, and consistent application of partnership principles.

A Key Debate: “Early Involvement.” The most significant tension involved what “inception” or “early involvement” meant. Labor argued for participation starting in the earliest stages associated with AI use—problem definition and assessment of potential solutions. Management raised questions about the feasibility of that approach, given confidentiality concerns and the decentralized nature of technology decision-making at KP. This debate ultimately shaped the agreement’s structure.

The Agreement: A Two-Pronged Approach

Rather than attempting to control all AI decisions through a single committee, the parties developed a model combining top-down and bottom-up strategies:

A Top-Down Strategy: The National AI Task Force

A jointly led task force with ten members (five labor, five management) will provide strategic oversight of enterprise-wide technology initiatives affecting multiple markets. The task force will:

- Share information about planned AI and technology rollouts
- Bring in subject matter experts and conduct learning sessions
- Guide communications, training, and education related to AI and other technologies
- Make recommendations for effective development, implementation, and workforce engagement
- Discuss concerns raised by either labor or management related to AI or other technologies, including topics touching HR functions like hiring and performance evaluation
- Evaluate the agreement’s effectiveness and recommend adjustments.

Critically, the task force will include senior leadership from both labor and operations, including those with direct decision-making authority over technology investments. The task force is charged with developing processes for early engagement of workers and labor representatives in technology decisions.

A Bottom-Up Strategy: Enhanced Unit-Based Teams (UBTs)

As a result of its long-standing labor-management partnership, Kaiser Permanente has more than 3,500 active “Unit-Based Teams” (UBTs). These teams consist of a combination of frontline KP workers, managers, and often providers such as doctors or dentists who work in the same department and meet regularly together to discuss ideas for improving the quality of patient services and the way work gets done. Recognizing that frontline workers possess essential expertise about operational challenges and opportunities, the new AI agreement strengthens existing UBTs to generate ideas for using AI and new technologies. This will involve:

- Providing innovation training so UBTs can identify problems and opportunities where technology solutions might help
- Improving the UBT Tracker, KP’s database of UBT projects, to make innovation ideas visible across regions
- Establishing clear pathways for ideas to scale and for bi-directional communication
- Creating intermediary structures (potentially using existing regional Labor Management Partnership Councils) to process and learn from distributed innovation.

The Preamble: Shared Values

In addition to detailing structures, the parties developed a preamble establishing three commitments:

1. **Relational Care.** The preamble states that “AI and technology shall harness the expertise of frontline workers, and complement clinical judgment and professional practice” while preserving the human touch in healthcare.
2. **Employee Voice.** Employees’ input will be incorporated throughout the AI and technology lifecycle—from problem identification through design, implementation, education, training, evaluation, and communication.
3. **Ethical Use.** The parties commit to avoiding discrimination, bias, and inequity; ensuring transparency; and maintaining accountability—including timely communication about any new monitoring technologies.

The Benefits of This Approach

The two-pronged strategy addresses both parties’ core interests:

- **For management:** A structured, predictable process for technology decisions, with clear roles and expected timelines, reduces delays and increases the likelihood of successful implementation. Input from frontline employees identifies real problems and design flaws early, before costly implementation failures occur.
- **For labor:** Workers and their representatives have genuine opportunity to influence technology decisions through both the task force (for system-wide initiatives) and UBTs (for identifying needs and generating ideas). Communication and transparency are built into the process, not added as an afterthought.
- **For patients and the organization:** Technology decisions informed by frontline expertise and designed with worker input are more likely to improve patient care and work experience while maintaining clinical judgment and the human elements of healthcare.

The model also reflects lessons from KP’s 29-year labor-management partnership, which has included successful initiatives like an Affordability and Competitiveness Task Force, adapted here for the pace and complexity of AI decisions.

Implementation Challenges and Next Steps

The agreement is a promising foundation, but significant work lies ahead:

Preparing for Bottom-Up Innovation. With over 3,500 active UBTs, KP and the Alliance must invest in:

- Basic AI training for UBT members
- Support for UBT coaches to help guide innovation
- “Matchmaking” capabilities connecting frontline concerns to technology solutions
- A cadre of “peer advisors”—union members trained in AI who can support colleagues in adopting these tools.

Launching the National Task Force. Critical decisions are needed regarding issues such as:

- Ensuring management members include executives who actually make technology investment decisions
- Establishing the task force’s working relationship with regional structures and UBTs
- Engaging AI vendors and developers to emphasize that KP seeks solutions to worker-identified problems, not just off-the-shelf tools.

Building Intermediate Structures. The agreement leaves to the task force the question of how regional and local bodies will process UBT suggestions and communicate learning across the system. Evidence from other organizations suggests that strong intermediate structures are essential for sustaining frontline engagement and generating a steady stream of high-quality ideas.

Addressing “Shadow Use.” KP and the Alliance should develop clear guidelines for worker experimentation with AI tools, addressing legitimate concerns about patient privacy and security while encouraging responsible exploration.

Broader Implications

While KP’s 29-year labor-management partnership history is unique among large U.S. employers, this agreement offers important lessons:

1. **Interest-based bargaining works.** Traditional position-based negotiations would not have produced this result. Genuine dialogue focused on interests, supported by problem-solving processes and subject matter expertise, created space for innovation.
2. **Labor solidarity matters.** The Alliance, representing multiple occupational groups and unions, presented a unified vision on AI. This solidarity was essential to achieving results no single union could accomplish alone.
3. **Worker voice requires structure and power.** Effective worker voice requires formal structures (like the task force and UBTs) that can provide repeated opportunities for influencing key decisions, as well as leverage that encourages openness to compromise.
4. **Frontline expertise is essential.** Technology decisions made without frontline input often create workflow disruptions, fail to address real problems, and generate resistance. Early and genuine frontline engagement improves outcomes.
5. **Movement toward worker-centered AI is growing.** Union agreements related to AI use are being reached across sectors—from hospitality (UNITE HERE with major hotel and casino operators) to entertainment (SAG-AFTRA and the Writers Guild of America, each with the Alliance of Motion Picture and Television Producers) and others. The labor movement is increasingly prioritizing AI issues in bargaining.
6. **Nonunion employers can learn from this model.** While unions provide the formal power that ensures worker voice is heard, the underlying principle—that frontline expertise should inform technology decisions—applies more broadly. Organizations without unions face the question: how will they ensure worker voice with sufficient weight to overcome traditional management resistance?

Conclusion

KP and the Alliance have accomplished something important: they have demonstrated that workers and employers can work together to shape how AI is developed and used in ways that benefit patients, employees, and the organization. The agreement is not a final answer but a beginning.

Success will depend on how well both parties invest in the structures created—particularly in building the task force’s credibility and ensuring UBTs have the resources and support to generate meaningful innovation. It will also depend on translating the agreement’s values and principles into consistent, daily practice across a large, decentralized organization.

The parties should be commended for maintaining their commitment to this agreement even as their economic negotiations proved contentious. Implementing these provisions effectively may serve as a path to rebuilding the trust and partnership that will be essential for all parties’ long-term success.



Note: The first draft of this Executive Summary was produced by a generative AI program (Claude) using MIT’s proprietary license to protect the confidentiality of the report. The authors edited the AI-generated draft to ensure it accurately summarizes the full report. The authors retain full responsibility for its contents.

Introduction

“

“The consequences of previous technological transitions for the welfare of workers and the strength of the middle class have depended not only on the nature and application of technologies but also on the frameworks and legal institutions that have shaped their design, adoption, and use and the distribution of economic surplus among owners, managers, and line workers. The middle-class prosperity that accompanied the second industrial revolution is owed in part to the success of labor unions and supporting legislation in negotiating for higher pay, reasonable hours, safe working conditions, and employment security.

In more recent decades, U.S. worker bargaining power has eroded as labor union representation has shrunk.... Many rights and norms are increasingly contestable in the era of AI.... The outcome for worker and societal welfare necessarily depends on the legal, regulatory, and bargaining regimes in place. Absent well-functioning institutions, the committee does not presume that market outcomes will be socially desirable ones.”

Source: National Academies of Sciences, Engineering, and Medicine, Artificial Intelligence and the Future of Work (Washington, DC: The National Academies Press, 2025), 7–8.

The above statement from a distinguished panel of the National Academies of Sciences, Engineering, and Medicine describes a growing consensus among experts about the biggest challenge facing the country in the use of artificial intelligence (AI): how to ensure that the workforce shares in the benefits of AI so that the U.S. avoids creating another generation of winners and losers from economic change.¹

The panel states the challenge well. Workers need to have a voice in shaping how AI is developed and used in their workplaces. Yet to date, no one has demonstrated how to meet this challenge.

This case study of Kaiser Permanente (KP) and the Alliance of Health Care Unions (the Alliance) serves as one example of how an employer and union took up this challenge as they negotiated a new collective bargaining agreement. We present the case study and the lessons we believe it offers as a path forward, both for Kaiser Permanente and its workforce, and for other employers, unions, and workers who face similar challenges.

This report documents and analyzes the negotiation process observed by the authors, based on fieldwork drawing on direct observation of an AI-focused subgroup during negotiations, with close attention to how issues related to AI and worker voice were discussed and addressed.

The report’s sections proceed in the following order. Following a brief description of the background to the negotiations, we provide descriptive and narrative accounts of the organizational context and the negotiation process. They are supported with details about how the authors gained access and their methods and the role they played as independent researchers. Then we synthesize the observations into a set of analytical findings that identify the patterns and processes that led to the development of an agreement on AI and other emerging technologies. The final sections discuss the implications of these findings for KP and the unions involved as well as for other practitioners and scholars.

Background

Kaiser Permanente (KP) is one of the largest not-for-profit healthcare plans in the U.S. It is an integrated healthcare organization that provides both insurance and healthcare services to 12.9 million members. KP has more than 240,000 employees, who work in hospitals and clinics located in California, Oregon, Washington, Colorado, Hawaii, Georgia, and the Mid-Atlantic states. The Alliance of Health Care Unions represents approximately 62,000 KP nurses, technicians, other professionals, and service employees.

KP has the most comprehensive and longest-lasting labor-management partnership in U.S. history. Our research group at the MIT Institute for Work and Employment Research (IWER) has had a long-standing relationship with both the labor and management parties to this partnership. Thomas Kochan, one of the authors of this report, has studied the evolution of the partnership since its inception in 1997, written about the partnership's phases,² and continued to follow partnership developments. KP and the Alliance agreed to let our research team observe and document the 2025 negotiations of the AI subgroup that was tasked with developing recommendations for how to work together on AI development and use.

Our commitment to the parties was that we would write a case study of their process and resulting agreement with two audiences in mind. First, we would prepare a report that we hope will be useful to KP and the Alliance, both in reflecting on their negotiations and in guiding the implementation of the terms of their agreement. Second, we want the case study to be useful to other employers, unions, and workers who are facing similar challenges in putting AI tools to work in their settings.

The labor-management partnership between the Alliance and KP offered an ideal opportunity for labor and management representatives to work together on all phases of AI decision-making. In previous bargaining and joint deliberations, the parties had demonstrated their willingness and ability to negotiate beyond the mandatory subjects of collective bargaining to address strategic issues facing the company and the workforce.³ They also have more than 20 years of experience in working together on the introduction and use of new electronic medical record technologies⁴ and thus are well-positioned to do so with respect to AI technologies. This case study allows us to flesh out what challenges were encountered as the parties used an interest-based, problem-solving approach to generate options for involving workers and their representatives in all phases of AI development and use. (See Figure 1 for a generic description of "interest-based bargaining" processes.⁵)

Figure 1. A Description of Interest-Based Bargaining (IBB)

"IBB is a collaborative approach to resolving labor and management disputes. Through the process, parties proactively identify durable solutions to outcomes at the bargaining table. Agreements are based on mutual and individual interests rather than positions. This approach emphasizes problem solving and enables mutual gain outcomes."

Source: Federal Mediation and Conciliation Service, "Interest-Based Bargaining: Joint Problem-Solving for Mutual Gain," <https://www.fmcs.gov/wp-content/uploads/2026/04/Interest-Based-Bargaining-2026.pdf>



Evolution and Key Features of the Labor Management Partnership at Kaiser Permanente

To put the 2025 negotiations in context, it may be useful to describe some of the history and key features of the KP Labor Management Partnership. The Partnership was born out of crisis in the mid-1990s. During the 1990s, KP had been experiencing intense competitive pressures and increasing labor unrest, including a number of strikes. The crisis came to a head in late 1995 when the Director of the Federal Mediation and Conciliation Service hosted a “summit meeting” of the KP CEO and leaders of the unions representing its employees to discuss what to do about their deteriorating labor-management relationships. That meeting resulted in a joint commitment to work together to build what turned into the largest and longest-lasting labor-management partnership in U.S. history.⁶

One of the partnership’s earliest major achievements was to negotiate a comprehensive Employment Income and Security Agreement (EISA) providing workers with protections against layoffs in times of restructuring or technological change. This proved especially helpful as KP invested heavily in electronic medical records technology. The EISA carries forward to today.

In 2005 the parties agreed to create what they call “Unit-Based Teams.” These teams consist of a combination of frontline KP workers, managers, and often providers such as doctors or dentists who work in the same department and meet regularly together to discuss ideas for improving the quality of patient services and the way work gets done. Today there are more than 3,500 UBTs active in KP workplaces.

The parties also negotiated and jointly administer labor-management training programs funded out of collective bargaining to provide workforce and career development opportunities for union members. The joint training program covering Alliance Union members is called the Ben Hudnall Memorial Trust, named to honor the memory of one of the partnership’s most ardent advocates for workforce training and development.

In 2018, the unions representing KP employees split into two large groups: the Alliance of Health Care Unions and the Coalition of Kaiser Permanente Unions. These two groups now negotiate separate agreements with KP.

Two of the key governing bodies in the Alliance-KP relationship that will be discussed in this report include the Labor Management Partnership Councils (LMP Councils), which are regional groups that coordinate and review partnership activities within the geographical regions in which KP operates, and the Joint KP Management and Alliance Executive Committee, which oversees and reviews partnership activities across the entire company. Also, in 2022 negotiations, KP and the Alliance created a national Affordability and Competitiveness Task Force to identify cost-saving and efficiency improvement opportunities that in turn help meet specific targets that trigger performance sharing bonuses across the workforce.

Since its first national-level negotiations in 2000, KP and its partnership unions have often used “interest-based bargaining” tools.⁷ (See Figure 1 for a definition of interest-based bargaining.) Subgroups are sometimes created to tackle specific issues such as work and family issues, attendance, workforce development, and so on. As will be discussed below, in these 2025 negotiations, the parties decided to create a subgroup devoted to addressing the growing challenges and opportunities associated with AI.

How the Parties Prepared for Negotiations


Before formal negotiations begin, labor and management organizations typically do a lot of preparatory work to identify issues that might be addressed, requesting and sometimes sharing information with each other, and ultimately developing priorities and processes to be used once the negotiations commence.

The Alliance began preparing for these negotiations in early 2024. Early on, the issue of what to do about the growing role of AI in healthcare became a topic of discussion. Similar discussions about AI were occurring among KP managers as they began preparing for bargaining in early 2025.

Our research team first learned about the possibility that AI would feature prominently in the upcoming negotiations when Hal Ruddick, the Executive Director of the Alliance, discussed the Alliance's interest in AI with us in June 2024 at the annual conference of the Labor and Employment Relations Association (LERA), an organization whose membership includes both management and labor leaders as well as scholars interested in labor relations. He knew of our interest in this topic and some of our prior work on the value of bringing workers' voice into AI decision-making. We discussed ways we might work with the Alliance and KP partners. Ruddick agreed to get back in touch after discussing whether or not this would be a central issue in the 2025 contract negotiations.

In the fall of 2024, Ruddick invited one of the authors of this report, Thomas Kochan, to present some of our work on AI in a Zoom meeting with the Alliance's Executive Board. We reviewed the current state of practice and developments in other union and management settings and offered two suggestions for how the Alliance and KP might address these issues in their upcoming negotiations. The first suggestion was to build on their practice in prior negotiations of creating a subgroup of union and management representatives to use interest-based bargaining methods to develop recommendations for members of the main bargaining table, which is called the Common Issues Committee, or CIC. The second was to consider use of the framework shown in Figure 2. This framework was developed by a number of us at MIT to highlight how worker voice can be integrated into all phases of AI strategy and implementation: problem definition; design of the technical and work process features; workforce education and training; and transition assistance and compensation for those whose jobs are affected.


Figure 2. *An Integrated Strategy for Technology Design and Implementation*
Incorporating Worker Voice Into Four Phases of Technology Design and Implementation



Credit: "Bringing Worker Voice Into Generative AI," by Thomas A. Kochan, Ben Armstrong, Julie Shah, Emilio J. Castilla, Ben Likis, and Martha E. Mangelsdorf. MIT Open Publishing Services (2024).

These ideas were well received by the Alliance leaders and led to a second invitation to make the same presentation at the December meeting of the Joint KP Management and Alliance Executive Committee. Subsequent to these meetings, the parties agreed to form a subgroup of union and management representatives that would be charged with developing recommendations for addressing these issues in negotiations. Figure 3 presents the charter drafted by the parties to guide the deliberations of the subgroup.

Figure 3. *The Charter of the 2025 AI and Technology Subgroup*



In February 2025, Kochan was asked to join a panel of KP and Alliance representatives and to co-host two hands-on workshops at a meeting of approximately 300 Alliance leaders and 80 KP leaders who came together to discuss their perspectives on the upcoming negotiations.

Observing and participating in this unique pre-bargaining conference provided an opportunity to get a preview of how AI would fit into the larger range of issues the parties would be taking up in negotiations. It also provided a window on the overall context in which these negotiations would take place. It is worth describing some of what took place in the two days of the conference.

It is important to note that holding a large joint labor and management meeting to discuss an upcoming negotiation is a relatively rare event in collective bargaining. Typically, each party holds its own pre-bargaining conferences/meetings to allow for frank conversations, internal debates and data sharing, and motivational speeches. But KP and the Alliance have a long history of also holding joint sessions prior to bargaining where the parties hear from each other about their expectations for negotiations. This conference also served as an opportunity for union delegates to talk privately, since the first day was open to both labor and management representatives and on the second day, the union representatives met on their own, without management participation.


The opening session of the first day featured speakers from the highest levels of KP and the Alliance. Two themes came through loud and clear in the comments from these leaders. One common theme was the hope that these negotiations would be different from the last time the parties negotiated an agreement, in 2021. Those negotiations were extremely contentious, mostly characterized by traditional arms-length bargaining, and went right down to the last minute, when an agreement was reached just short of a strike. Both union and management speakers, including Ruddick and KP CEO Greg Adams, noted that this time they hoped to use this bargaining to address critical challenges.

The second theme in the remarks differed somewhat between KP and Alliance speakers. Alliance leaders stressed the need for the wage settlement to make up for ground lost to inflation, particularly because Alliance unions' 2021 negotiations had yielded smaller wage increases than other KP union negotiations in 2023. KP leaders recognized these union concerns but stressed the much more difficult economic conditions facing KP and other healthcare providers in 2025 and in the near-term future. This difference in economic perspectives—with the union looking back in time and focused on the economic stresses the membership experienced and the employer looking forward to difficult economic times to come—is a typical indicator that wage bargaining would be difficult. Indeed, differences in time perspectives—one party looking at prior lost ground and the other focused on future challenges—is often a good predictor of a strike in labor negotiations.⁸

The joint conference also featured a panel discussion that included Jane Carter, Director of Research, Public Policy and Regulatory Affairs at UNAC/UNHP, a large union of nurses and other healthcare professionals in Southern California and Hawaii; Bryan Wynn, International Vice President and National Bargaining Director at the United Food and Commercial Workers International Union; Sam Glick, Executive Vice President of Enterprise Strategy and Business Development at KP; and Kochan, a member of our research team. Some of the materials presented in the prior briefings were again presented to this larger set of labor and management leaders. Once again, the framework shown in Figure 2 was presented and was accompanied by Glick's presentation of KP's "PROTECT" framework, which is shown in Figure 4 and outlines KP's principles for using AI. Carter and Flynn discussed how their unions were developing AI strategies and looked forward to exploring these issues with KP colleagues in the upcoming negotiations. That session helped stimulate thinking and build support from both labor and management participants for addressing these issues in a problem-solving/interest-based fashion.

Additionally, there were two workshops in which conference participants discussed their experiences with AI to date, their concerns, and their hopes for using AI on their jobs. Carter and Kochan facilitated these workshops. The discussions were lively and informative. The key lesson Carter and Kochan took away from the discussions was that there was wide variation in the participants' views on and experience with AI. Some clearly feared that AI will take away their jobs. Others reported how they were already using aspects of AI to support their work. Still others were curious about how they could gain a voice in shaping the use of AI tools that would likely be coming to their jobs in the future. The common theme shared by workshop participants was that the upcoming negotiations presented a great opportunity to clarify and help shape how the workforce might engage on AI issues going forward.

Figure 4. Kaiser Permanente’s PROTECT Principles



Invitation to Observe AI Deliberations in Negotiations

After these sessions but before negotiations began, KP and Alliance leaders agreed to allow our research team to sit in on and observe the work of the AI subgroup in the negotiations process. The terms of our mutually agreed-upon role were outlined in an April 9, 2025 email from Jim Pruitt, Vice President, Labor Management Partnership and Labor Relations for the Permanente Federation:

- “You and your team are approved to observe the Alliance-KP National Bargaining with these understandings:
- 1. Your focus is the AI Subgroup.
 - 2. You are observing the process and drawing conclusions but not advising the negotiators.
 - 3. You will not comment on the negotiations until there is a settlement.
 - 4. Before you publish you will give Kaiser Permanente and the Alliance of Healthcare Unions an opportunity to review a draft of your work and to suggest input.
 - 5. You and members of your team are welcome to sit in the plenary sessions.
 - 6. You and members of your team are welcome to sit in caucuses of the AI group only with the approval of the management or union leader of the caucus. In full transparency, at this point management is disinclined to approve this.”

Our role thus shifted from one of providing advice on how to address these issues to one of independent researchers observing the parties’ negotiations in order to document the process and results and to prepare this case study. The only deviation from the observation role was that, at the mutual request of the parties, Kochan repeated the presentation on AI made at the joint leadership session for the benefit of those labor and management members of the subgroup who were not at any of the prior presentations.

Negotiations for the new contract opened on May 5, 2025. Between then and August 22, 2025, the subgroup on AI met six times, with each meeting extending over the course of two to three days. The subgroup consisted of seventeen representatives from the different Alliance unions and seventeen management representatives. An additional twenty union representatives and a small number of management representatives attended the sessions as observers. Two facilitators were present at each session to support the interest-based problem-solving processes.

The authors collected data primarily by taking field notes on observations during the negotiation's plenary sessions, the AI subgroup's meetings and small group discussions, and labor caucuses, though not during any management caucuses. These field notes were supported by informal discussions with union and management participants during breaks, meals, and communications between sessions. In addition to the subgroup participants, we also talked informally with the labor and management observers who sat in on negotiations and contributed to the caucus discussions. The idea of inviting local union members to observe was to provide the various local unions opportunities for members to see how the parties conducted the interest-based negotiations.

At least two authors were present at each session that the research team observed, allowing for overlapping note-taking. This supported the accuracy of detailed observations and enabled coverage of simultaneous small group discussions during subgroup meetings. The authors met at the end of each day to review the day's events, compare interpretations of various interactions and dynamics, and plan subsequent observations and informal discussions. Following each negotiation session, at least one author wrote analytical memos that explored emerging themes and preliminary interpretations of the negotiation process.

The authors attended the sessions as invited, independent observers, not as advisors or decision-makers. We all have prior experiences studying worker voice and labor-management partnerships, which provided multiple perspectives on negotiation dynamics and sensemaking in settings with pluralistic interests and values. In line with rigorous qualitative research practice, we approached this work with ongoing attention to how our position as researchers could shape our interpretations, so we aimed to hold each other accountable not to advance any particular position or solution. We focused on documenting how participants made sense of emerging technologies and partnership structures and on representing the reasoning and experiences they expressed regarding their hopes and concerns about AI and other new technologies.



How the AI Subgroup Approached Its Work

To prepare for its work, members of the AI subgroup spent time learning about principles and cases about AI implementation through presentations by subject-matter experts and readings on AI in healthcare. The subgroup also observed three virtual and two in-person presentations from technological, clinical, and organizational subject-matter experts who described different models for implementing AI and managing technology changes more generally.

- Daniel Yang, MD (VP of AI and Emerging Technologies at KP): Dr. Yang's virtual presentation provided a high-level overview of different types of algorithmic systems (i.e., classification, predictive, and generative) and explained how these can affect patient care, risk, and safety. He emphasized the importance of defining the problem before selecting a solution and the central role of trust in supporting implementation.
- Professor Eva Selenko (Professor in Work Psychology, Loughborough Business School): Professor Selenko presented virtually from the UK about the psychological impact of working with AI, including the risk of reduced autonomy, meaning, and motivation, as well as potential shifts in professional or occupational identity because AI technologies are able to complete some tasks.
- Kirk Tamaddon, MD (Area Medical Director and Chief of Staff at the Kaiser Permanente West Los Angeles Medical Center): Through a virtual presentation, the subgroup learned about the introduction of the da Vinci surgical robot system from Dr. Tamaddon, a pioneer in minimally invasive and robotic surgery. The presentation underscored that successful adoption required clinicians to be directly involved as experts in evaluating the technology's feasibility and workflow integration, with discussions of involving labor partners, too.
- Thomas Kochan (Massachusetts Institute of Technology): A member of our research team presented (in person) a framework for worker involvement in technology decision-making, previously summarized in this report.
- Kyle Arnone (Senior Manager, Center for Collective Bargaining at the American Federation of Teachers): Arnone described (in person) how the American Federation of Teachers established a collaborative partnership with several AI companies, bringing worker-centered concerns into the development and design of new AI tools, as well as into the training employees received about using the new tools.

In addition to these presentations, the subgroup circulated and discussed in small groups two practice-oriented articles that offered concrete cases of AI implementation in healthcare.

- "AI on the Front Lines" discussed the integration of AI decision support tools at Duke Health.⁹ This article highlighted how frontline resistance can stall adoption if workers see low benefit, added labor, or reduced autonomy. It emphasized engaging frontline workers early to identify pain points, working directly with vendors to adjust tools based on user feedback, and protecting professional autonomy within clinical workflows.
- "Scaling Automation: Two Proven Paths to Success" reviewed practices at two hospital systems, Mass General Brigham and Mount Sinai Health System.¹⁰ This article described how both systems relied on automation scorecards to evaluate potential automation projects. The scorecards helped identify tasks suitable for automation by assessing automation fit, level of effort, and value created, offering a structured way to prioritize new technology use cases.

Throughout the sessions, subgroup members also made selective use of generative AI tools to support their conversations and idea generation, while being careful not to include any confidential organizational information in prompts. In small-group activities, we observed participants using OpenAI's ChatGPT and Microsoft Copilot to summarize and categorize the interests of labor and management and to suggest potential principles for responsible AI use.

One small group even asked the tool to generate a short poem about the opportunities and challenges of AI in healthcare; this may well be the first AI-generated poem to emerge, even in this whimsical way, in a collective bargaining negotiation! (See Figure 5.) Small groups using the generative AI tools did not just accept outputs as finished products or recommended positions. The generated ideas served as starting points to discuss and refine together. This process allowed the subgroup to surface different perspectives during discussions and to clarify which ideas should guide next steps. Debating the AI-generated results helped the group align and move forward.

Figure 5. *An AI-Generated Poem Prompted and Discussed in a Small Group*

“Trust the Tech, Fear the Unknown”

*They say AI is coming—
but what’s coming for me?
Surveillance whispers behind glass screens,
tracking steps I never knew were seen.
Watching. Measuring. Comparing.
Will my worth be algorithmically declared?*

*I want control—
of my job, my hands, my name.
I want security, not speculation,
not data dictating my promotion or erasure.*

*They say trust the process—
but trust is earned, not automated.
It’s built in conversations,
not calculations.*

*So talk to me.
Give me words I can hold,
transparency I can taste.
Teach me the ethics of tomorrow,
before it redefines today.*

*Let AI elevate, not replace.
Let it serve patients, not overshadow their faces.
Let labor lead the leap,
and keep humanity in the code.*

*If machines shape the future,
then we shape the machines.
And if trust is the key,
then let’s build it—not bury it.*



How Participants Understood the Problem and Articulated Their Interests

The AI subgroup's discussions about AI and emerging technologies were motivated by the claim that these tools have the potential to reshape healthcare delivery, organizational operations, and the daily experience of frontline employees. Across multiple sessions, participants articulated specific interests that we review here. In addition, the participants' dialogues, debates, and strategic silences sometimes pointed to other interests and concerns that were less explicitly stated and that the research team had to infer from what we observed.

The operating assumption throughout the sessions was that progress on figuring out how to partner on AI and emerging technologies in practice would depend on clarifying the motivations, concerns, and needs of all parties. Facilitators frequently used examples to distinguish between positions ("who gets the orange") and interests ("why the orange is needed or valued by different parties"), emphasizing that a bargaining conversation grounded in understanding each other's underlying interests would be more likely to solve real problems and achieve both shared and distinctive goals.

While there were clear areas of alignment in interests—including commitments to patient care, operational effectiveness, and a healthy labor-management partnership—we repeatedly noted that the parties had different interpretations of core concepts such as early involvement, expertise, and organizational agility.

We summarize the expressed and emergent interests around AI and new technologies in four broad categories:

1. **Patient and mission-centered interests**, focusing on patient care quality and patient trust.
2. **Job quality and job security interests**, including keeping the "human touch," meaningful work, autonomy, fairness, and concerns about job security.
3. **Innovation interests**, including organizational agility and the need to move quickly with new technologies to stay competitive.
4. **Systemic or process-centered interests**, such as the desire for effective communication, timely implementation and the avoidance of roadblocks, and the appropriate structures to support ongoing partnership regarding AI and other new technologies.

Participants expressed hope that AI and other new technologies may enable improvements in all four areas, but a pervasive theme in the subgroup discussion was that such benefits were more likely if workers' insights were meaningfully integrated into key decisions. That working assumption was repeatedly stated by labor members but also referenced by management participants. A key topic of discussion, therefore, was what constitutes appropriate input from frontline workers and their representatives, and how such input should be balanced with organizational needs for speed and flexibility.



An illustration at the top of the page shows the silhouettes and partial faces of several people of different ethnicities and ages, suggesting a diverse group in a meeting or discussion. The colors used for the silhouettes are green, orange, purple, yellow, and red.

Patient and Mission-Centered Interests

Improving patient care and experience as the starting point. An early presentation, setting up the subgroup's charge, emphasized the potential for AI to advance core patient-care goals and signaled the expected centrality of those interests in the discussions. This presentation highlighted anticipated benefits for patients, such as enhanced diagnostic accuracy, more personalized treatment plans, streamlined processes and gains in efficiency, and the potential for reduced costs, which would be helpful in sustaining the organization. In later discussions, those goals were framed as extensions of the organization's mission to deliver high-quality care. Several times, other members of the subgroup reinforced this orientation to high-quality care by focusing on how technology might support the "human" aspects of care, as discussed below.

Patient trust and data stewardship. Participants cited fears among patients—particularly in communities with historical reasons for skepticism, such as Black communities—that their data might be misused, mishandled, or insufficiently protected. The integrity of patient data practices was thus seen as essential not only to ethical practice but also to maintaining long-term patient loyalty and organizational credibility. Participants frequently linked their own uncertainty about AI to what they believed patients would feel: confusion, fear, and worry about loss of human connection. Some suggested that if workers themselves felt secure and informed, they would be better positioned to help patients navigate technological change.

Meaningful work and the "human touch." Members of the subgroup repeatedly described the importance of maintaining meaningful work, with references to autonomy, discretion, and the ability to connect with patients or colleagues in ways that felt professionally and personally fulfilling. Many expressed hope that AI would ease burdensome or repetitive tasks, reduce burnout, and create opportunities for new skills or professional growth. AI was described as a possible strategy to free clinicians or other staff from documentation burdens, allowing more time for working at the "top of one's license" (i.e., doing work that uses one's full skills and training) and for patient interactions focused on connection, explanation, or reassurance. Notably, though, the discussions of maintaining the human touch and using AI to improve the work experience were often brief, stated as an interest but not elaborated in detail. We found the explicit discussions of "the human touch" interesting, since it reveals that health care is being conceptualized and discussed as an interaction that may not necessarily involve a human delivering the care. Maintaining and supporting "the human touch" only becomes an interest when the absence of human interactions is conceptually feasible. Additionally, while early discussions of the importance of AI for healthcare stressed reducing burnout and addressing labor shortages, members of the subgroup rarely spoke explicitly or in detail about how AI would practically advance those goals.

In parallel with the hope that AI would support meaningful and valued work, subgroup members also voiced concerns that poorly designed tools could undermine job quality by introducing rigid workflows, interfering with professional judgment, or reducing roles to more routinized tasks. When these concerns surfaced, they were often tied to past experiences in which new technologies or systems added administrative burdens or were not aligned with frontline needs. We saw that these stories were usually offered by labor members of the subgroup and often received with great interest by management participants, perhaps because the management representatives to the subgroup did not include many managers immersed in clinical operations. This meant stories from the frontlines and from providers were viewed as useful, and sometimes surprising.

While early discussions of the need to keep the human touch and keep a human in the loop were brief and seemed to point to a common interest, there were more questions and some challenges getting alignment about what those stated goals meant as the negotiations unfolded. These broad goals reveal an interest in protecting human decision-making, both because overreliance on AI systems is risky and degrades skills and because employees value a sense of autonomy and the ability to make decisions based on their expertise. Management members of the subgroup repeatedly noted that licensed providers are legally required to make decisions and are legally accountable for those decisions. The implication was that professional licensing already provides an adequate backstop for ensuring human authority to “override” algorithmic prescriptive recommendations (i.e., guidance from AI systems). Indeed, the term “override” became contentious in the later discussions. Labor members raised questions about how the expertise of staff who do not have professional licenses would also be valued, since their unions represent many different occupational groups. Labor participants’ comments suggest a position that grounded occupational expertise may lead those workers to make legitimate decisions that differ from AI recommendations and that their decision authority should be protected.

Job Quality and Job Security Interests

Job security, employment stability, and fairness. An interest in job security was implicit throughout the discussions; it was often referenced in passing but not laid out in detail. For example, a small group reporting out to the subgroup noted an interest in “overall job security, making sure we keep people” without feeling that they needed to explain that further. The brief references to this core interest reflect the fact that employment impacts and related adjustments had already been addressed and covered by the comprehensive Employment and Income Security provisions in the existing contract and therefore need not be addressed by the subgroup. We surmise that this limited discussion also reflects all parties’ desire to be viewed as realistic about and open to AI and other new technologies, despite possible changes in work processes and some job loss.

An expressed interest that was articulated by labor members as negotiations unfolded was ensuring bargaining unit integrity, i.e., recognizing that jobs might be restructured substantially enough that job titles and occupations changed but those entering the newly changed jobs should be represented by the same union, or perhaps another union representing that occupation.

Labor members also voiced a desire for greater visibility into how AI might influence hiring, promotion, and expectations for productivity or performance metrics, both currently and in the future. Some expressed concerns that algorithmic tools could shape employment decisions without sufficient transparency or understanding of underlying data. These discussions were consistent with broader concern about losing control over one’s work, including the possibility that AI systems could monitor work processes and prompt corrective or disciplinary actions without the affected employees’ knowledge or understanding of those systems.

Discussions of fear often came up in conjunction with concerns about potential surveillance, loss of autonomy, job insecurity, and a perceived lack of transparency about existing technologies. Labor representatives repeatedly asserted that providing more information about how AI is used in internal HR systems would help to “reduce fear of AI,” which would minimize frontline workers’ concerns about AI tools and support the implementation of new systems or features. Additionally, labor representatives stressed that their personal relationships with frontline workers were an important resource for getting past workers’ fears; if labor leaders were ready to offer reassurances about AI governance and about the AI strategy more generally, then workers would be more open. Management participants seemed to accept this claim and recognize that labor representatives would play a pivotal role in frontline workers’ responses to AI and new technologies.

Training opportunities. There was a stated interest in robust training and retraining systems. While existing programs (e.g., the Ben Hudnall Memorial Trust training fund for Alliance members) were viewed as valuable for long-term development of skills and career advancement, some members of the subgroup saw them as insufficient for helping workers adapt quickly to technological change or transition into alternative roles if displaced. Several people stressed that training must be timely, practical, and tailored to the evolving uses of AI across the organization. However, the details of how to pursue this interest received less attention in the subgroup, perhaps because the plans for developing this training and retraining system were eventually passed on to a proposed national task force.

One perspective that we found notable for its aspirational vision came from a labor member who was a licensed clinician. Because a new technology that affected their own work had previously prompted significant “professional growth,” this person proposed that AI could be framed as an “inspirational and positive” opportunity. Reflecting on the skill development experienced during this energizing period of learning tied to a new technology, they shared that “I would hope that implementing new tech can make people feel that way. I want that for my colleagues.”

Innovation Interests

Speed, agility, and competitiveness. Management members of the subgroup frequently referenced the need for agility in evaluating, purchasing, and implementing new technologies. The pace of AI development, combined with competitive pressures in the healthcare industry, were described as requiring faster decision cycles and streamlined processes. Some noted that delays in implementation could undermine patient care, frustrate clinicians and other staff, or place the organization at a strategic disadvantage. This was the most directly and frequently stated interest by management members.

At the same time, both labor and management members acknowledged that moving too quickly and without adequate worker involvement could create inefficiencies, produce unintended harms, or require costly rework. Labor representatives repeatedly argued that early worker engagement would actually accelerate implementation by ensuring that potential barriers or design flaws were identified early. Members pointed to specific experiences where tools were purchased or deployed before workers had an opportunity to provide input, or cases in which input was invited but not incorporated into the new tool as eventually deployed. Labor members explicitly said that such experiences not only delay effective implementation but also reduce workers’ confidence in the partnership.

Balancing agility with partnership commitments. A recurring challenge was how to reconcile the interest in speed with the interest in having worker involvement early on. Management members initially raised concerns about how to move quickly given pragmatic questions about the capacity of labor representatives with busy schedules to weigh in effectively on complex and unfamiliar technologies, as well as concerns that confidentiality requirements and the currently decentralized nature of decision-making would mean labor input would be fairly slow. Different technologies require different types of expertise, and identifying the appropriate representatives at each stage was described as complex.

Management representatives emphasized that overly complex structures or unclear expectations could slow decision-making. Employees, conversely, argued that insufficient early involvement of labor produced the very roadblocks that management sought to avoid. We believe that, early on at least, there was a preferred management position of avoiding the establishment of additional committees or consultative structures; that preference likely reflected an interest about limited staff time and resources, as well as agility. This position was referenced in small group discussions and informal conversations, but it was not elaborated upon or discussed directly in the broad subgroup deliberations.

Despite these differences, there was general agreement that more effective and predictable processes—combined with shared understanding of expectations—could reduce delays and improve outcomes. Eventually, the national task force was proposed as a strategy to balance moving quickly with appropriate input.

Process-Centered Interests

Clarity, transparency, and consistent structures. An interest in clear and consistent structures that enable effective communication, support timely implementation, and reinforce partnership expectations was articulated by both labor and management members of the AI subgroup. Many noted that existing processes for technology decision-making were difficult to navigate, and uncertainty was expressed about who made decisions, how information flowed, and where input could be provided.

Labor members repeatedly voiced requests for comprehensive inventories of existing AI tools or how AI is incorporated in current technologies. Management members repeatedly responded with the observation that such an undertaking may not be feasible. We see this management response as highlighting the tension between the desire for transparency and clarity, expressed most often by labor partners, and the reality of decentralized or informal decision pathways. Our assessment is that some of this ambiguity about current uses and decisions reflects KP's very decentralized system (with different markets and subunits making decisions and purchases separately), but this may have also reflected the limited representation of IT professionals in the bargaining subgroup.

Communication, feedback loops, and accountability. We heard frequent references to communication challenges. As part of the broader discussion of the balancing of speed and input, subgroup members shared prior instances in which changes were not communicated effectively, leading to workflow disruptions or misunderstandings about expectations. Many argued that communication strategies should be co-developed by labor and management and should include timely feedback loops so that frontline experience can inform adjustments.

Some labor members also stressed the need for accountability when agreed-upon communication or involvement processes are not followed. They noted that partnership principles sometimes failed to reach operational levels, resulting in inconsistent implementation and eroding trust. Several labor representatives emphasized that AI initiatives would advance only “at the speed of trust,” using a phrase from Dr. Yang's opening presentation. Many described trust as an outcome of transparent processes, early involvement, and reliable communication.

Differing Understandings of “Inception”

Both labor and management representatives articulated a shared interest in workers' involvement in AI and technology decisions. A management member of the AI subgroup stated explicitly that there is a common interest in having a “diversity of thought when it comes to technology and innovation.” Several labor representatives claimed that such input and involvement is already required in existing agreements and is a foundational expectation in the labor-management partnership.

However, a central area of debate was what it means for workers to be involved “at inception” of new technologies or even to have “early involvement.” For frontline employees and their representatives, inception generally referred to the earliest stages of problem definition, identification of needs, and assessment of potential approaches. They argued that problems should be defined collaboratively, and that frontline expertise is essential to determining whether a proposed tool addresses real operational challenges, i.e., in phases 1 and 2 shown in Figure 2.

Management members frequently agreed in principle that early involvement can be beneficial but also raised questions about feasibility. In some cases, the earliest realistic stage for worker involvement was described as the beginning of implementation planning, rather than during ideation or during the development or evaluation of products. This seems to reflect the previous point that technology decisions are currently quite decentralized, require specialized technical knowledge, and are not clear to all parties, including, in some instances, management members of the subgroup.

Our impression is that there are two additional concerns about involvement from “inception” or “early involvement” from management’s perspective; these concerns were mentioned only briefly and so we offer this description tentatively. First, an expectation of early frontline input (i.e., in problem definition and identification of potential technologies to either purchase or build) implies that different technology decisions and purchases will be made across markets and perhaps across units in the same organization, which may complicate implementation efforts. A management interest in efficiency and consistency was articulated, tied specifically to “enterprise-level” implementation. This concern may reflect management’s recognition that successful, full implementation of new tools usually requires integration into existing systems (e.g., electronic health records) and workflows, although we did not hear that rationale offered explicitly. Second, a management member noted that an explicit commitment in the contract language to involving workers at “inception” could prompt frustrations and grievances because it is hard to define. He noted that ideas and pilots may arise from many directions, meaning it can be difficult to catch all initiatives early and to determine whether workers or labor representatives had been pulled in early enough to fulfill such contract language. Arguably, a bargained commitment to involving workers “at inception” would be hard to uphold and might result in many grievances. This practical concern, though, could function to limit early involvement and so make it difficult to achieve the stated interest in using AI to solve frontline workers’ problems and support the shared goal of improving care quality and access.

The meaning and distribution of expertise. Underlying these debates were differing views about expertise. Some members, particularly those on the management side who have been involved in the procurement of technology services and systems, emphasized the specialized knowledge required to evaluate certain technologies, suggesting that certain professional roles must have the primary say in selecting the appropriate tools and in negotiating with vendors.

Frontline employees, however, stressed that their experiential knowledge is essential for understanding workflow implications, patient needs, and practical feasibility. Several labor members described situations where frontline input was solicited only after key decisions had been made, which contributed to them feeling that their expertise was undervalued. Others noted that without clear channels for frontline ideas, opportunities for innovation may be lost. These points were taken up in the eventual proposal for KP’s Unit-Based Teams (UBTs) to surface concerns that might be addressed with new technologies and to track whether AI projects addressed these concerns.

Additionally, we noted that some subgroup members explicitly valued institutional expertise, i.e., they appreciated deep familiarity with the partnership agreements and processes such as UBTs. This institutional expertise was seen, by some at least, as an important resource for problem solving and for seeing ways to build on previous partnership successes and existing structures.

Toward a shared framework. Despite differences, participants acknowledged that both clinical or technical expertise and frontline insights are necessary. The challenge, therefore, is not whether workers should be involved, but how to structure involvement so that it is timely, meaningful, and compatible with the need for agility. Eventually, the subgroup proposed a national task force as a structure capable of integrating worker voice into rapidly evolving technology decisions, while also charging UBTs with identifying needs that might be addressed with new technologies.

Developing Both Top-Down and Bottom-Up Strategies

Having spent time articulating and understanding the various interests of the parties, the subgroup moved on to developing and evaluating various proposals to address the primary interests. The subgroup ultimately aligned on a model with both top-down and bottom-up strategies. In Figure 6, we summarize key components of the agreement. (The full tentative agreements submitted to and approved by the Common Issues Committee can be found in the Appendix to this report.)

Top-Down Task Force

The subgroup agreed upon an Alliance-KP Task Force on Technology and Artificial Intelligence with equal representation from labor and management. This body will serve as a shared venue to receive and distribute information on emerging technologies across the system and to make recommendations related to development, implementation, training, and communication.

The task force will serve as an enterprise-level convening body on emerging technologies that affect multiple markets; it will allow for early engagement with labor representatives as well as with members of management representing different functions and expertise. The issues the task force will address are described in the agreement as “national initiatives” but were discussed as any technology change potentially affecting multiple regions, as well as learning and communications strategies that have implications across KP. The task force will receive information and hold strategic discussions about planned and prospective AI and technology initiatives, bring in relevant subject-matter experts as needed, and develop recommendations related to development, implementation, training, and communications. Its responsibilities are most clearly defined as they pertain to training and communications, where, according to the agreement, it is expected to “guide” activities and identify needs, then work with existing bodies and staff. There is less specificity regarding how the task force will engage with earlier stages of the technology lifecycle, such as development and implementation planning, though the task force is asked to develop processes to address those issues. The task force is also charged with ensuring that a timely evaluation is carried out to assess the effects of the AI agreement in practice and suggest changes in processes.

As is to be expected with a new initiative, the agreement emphasizes goals and the core structure of this body rather than specific practices or decisions. The task force’s expected outcomes include increased communications, collaboration, and improved implementation of chosen technology solutions. In the agreement, there is also an expectation that the task force will “promote workforce engagement in development, adoption, and rollout” of AI and other technologies and that the technologies and tools adopted will orient to the goals of “improved patient care, efficient operations, and improved work experience.”

Bottom-Up UBT Enhancements

The idea that AI technologies are particularly well-suited to “bottom-up” innovations was introduced in the presentations made to labor and management representatives prior to the start of negotiations. This idea featured prominently in a 2024 MIT white paper that outlined some initial ideas on how to bring workers’ voices into AI development and use. That paper stressed that, unlike many other technologies such as robotics or electronic medical records that required large capital investments before reaching the workforce, AI tools are accessible via laptops, cell phones, and other devices widely available to workers.¹¹ As mentioned above, many of the KP workers we talked with in pre-negotiations workshops already had experiences using AI tools in their personal and work roles. Labor and management representatives on the subgroup recognized this from the start of their deliberations and therefore focused less on questions of whether or not to encourage and support this approach and more on how to fit it into existing workplace practices.

As a bottom-up strategy, the subgroup agreed to strengthen KP's existing Unit-Based Teams (UBT) structures so that frontline worker experiences and insights can more directly drive innovation and inform technology adoption. This approach involves innovation training so UBTs are well suited to identifying issues where AI and other technologies could perhaps enhance workflow or patient experience on the frontlines. With innovation happening locally across the enterprise, there is the hope for "increased engagement and empowerment of frontline staff, accelerated identification and scaling of innovative solutions," resulting in "improved service delivery and operational efficiencies." (Those quotes come directly from the agreement.)

This agreement calls for regional Labor Management Partnership Councils or some other body specified by the national Task Force to examine and improve communication across UBTs and regions and to enhance the current UBT Tracker, a web-based application used to collect information on local UBT projects; the aim is for UBT-driven innovation to travel up and across the organization in a timely manner. While these bodies are expected to serve an intermediary function among UBTs, the agreement does not fully specify how coordination will occur between the national task force and frontline UBT activities. This creates a degree of uncertainty around how top-down initiatives will be translated into local practices and how frontline insights will inform system-level technology strategy and decisions.

Figure 6. *Key Components of the Agreement*

	Task Force	Enhancing Unit-Based Teams (UBTs)
Objectives and Primary Functions	• Share information about planned innovation, technology, and artificial intelligence rollouts • Learn together about how artificial intelligence and other new technologies are impacting KP and healthcare more broadly • Guide communication, education, and training related to artificial intelligence and other technologies • Evaluate effectiveness of the agreement	• Empower frontline workers to identify concerns for which AI and other new technologies might be useful • Enhance supporting mechanisms for innovation within UBTs • Establish a clear pathway for innovative ideas to scale across business units and to encourage bi-directional communication, so frontline ideas and potential use cases are heard by task force and others
Participants	Ten members (five labor and five management), including senior leadership from labor and operations, plus executive leadership with experience in technology	Frontline workers and managers participating in UBTs. There was discussion of whether regional Labor Management Partnership Councils or another body would play a coordinating role, but that was not explicitly resolved.
Primary Functions	• Discuss AI projects in their development phase • Identify and recommend training needs • Be a connective hub for AI and other technology topics • Strategize on communications regarding AI and other technologies	• Train and communicate with employees about AI topics • Identify needs, potential use cases, and ideas for using existing and newly designed technologies • Improve the existing UBT Tracker system so it can capture technology innovation opportunities across UBTs and regions

The Agreement's Preamble

After the subgroup agreed upon the components of the top-down and bottom-up model, they developed what they called a preamble for the agreement to clarify the shared values and objectives—what one subgroup member referred to as the “theory” of the agreement. (See the Appendix to this report. The introductory paragraphs of the Appendix constitute the preamble to the agreement.) The preamble centers on three commitments: relational care, employee voice, and the ethical use of AI and other technologies. The statement on relational care clarifies that these technologies should support patient-centered care and employee experience while preserving the human touch, clinical judgment, and professional practice. The statement on employee voice calls for employee input throughout the technology lifecycle. The statement on ethical use highlights transparency, accountability, and the need to avoid bias, including timely communication with the workforce when technology involves monitoring.

In developing the preamble, the subgroup members drew primarily from a running list of labor interests that the subgroup could not incorporate into the task force or UBT enhancement agreements. The relational care language drew from discussions about maintaining the human interaction that patients prefer. Additionally, this section represented the resolution of “override” discussions by stating that “AI and technology shall harness the expertise of frontline workers, and complement clinical judgment and professional practice.” (The phrase “professional practice” was included to incorporate the expertise of non-clinical employees.) The employee voice section was built on frameworks that the subgroup discussed: the integrated strategy for including employee voice in the technology lifecycle shown in Figure 2, and KP’s PROTECT framework shown in Figure 4. Labor members referenced these frameworks as a way to clarify points of engagement given debates around “inception.”

The ethical use section of the preamble emerged from concerns about AI applications in HR functions such as hiring and performance monitoring. While specific responses to these concerns were not incorporated in detail, the preamble establishes expectations for timely communication, “per local practice,” when AI technologies related to monitoring productivity or performance are implemented. Additionally, the preamble notes that topics of concern to either labor or management, including the use of AI in HR functions, can be taken up at local levels and speaks to broad commitments to avoiding bias, discrimination, and inequity, and maintaining accountability and transparency.

The preamble captures the subgroup’s shared values but does not establish clear, enforceable guidelines or guardrails. We noted that key implementation questions are left to be worked out through the Task Force and UBT structures, and guidance on more contested topics remains ambiguous.

How Did They Get There? Resources and Key Activities

The AI and new technologies agreements were shaped by two major features of the subgroup’s process. First, the AI subgroup members drew on insights from the subject matter expert presentations they heard and the readings on AI in healthcare. Second, they drew heavily on successful existing partnership structures, particularly the Affordability and Competitiveness Task Force that KP and the Alliance developed in 2021. This task force offered a practical template for how a joint body can sponsor system-wide work, applied here to AI and emerging technologies.

The presentations and readings prompted the subgroup to reflect on how technology decisions move up and down the organization, and why both top-down structure and bottom-up expertise are needed to achieve responsible, effective implementation. Subgroup members were most drawn to how these cases framed frontline workers as essential partners in the design and deployment of technology, instead of just being downstream recipients.

In addition to information gained about best practices from experts and external cases, the opening plenary sessions before negotiations began underscored examples of successful partnership inside KP, such as UBTs and the Affordability and Competitiveness Task Force, (often referred to by participants simply as the Affordability Task Force, or the ATF). Presenters emphasized how these structures have historically strengthened trust, improved performance, and empowered workers to shape operational and clinical decisions. By highlighting these successes at the outset, partnership structures remained front-of-mind for subgroup members.

The ATF stood out as a model of an innovative and collaborative body capable of tackling system-wide challenges. Members of the AI subgroup highlighted that the ATF succeeded for several reasons. It combined regular, substantial meetings with real resourcing, a group of high-level representatives, supported by a professional facilitator, and a clear charter grounded in the National Agreement. They emphasized the ATF's ability to share best practices across the system. These features convinced the subgroup that an AI Task Force needed similar elements: defined membership that included both high-level executives with direct responsibilities for AI and high-level union leaders, a well-resourced body with a cadence of meetings relevant for AI, the explicit charge to guide multi-market technology efforts, and regular communication with regional structures and local UBTs.

Other Proposals Discussed

Importantly, along the way, the subgroup considered other ways to meet the deliverables of their charter. The following alternative proposals were eventually interpreted as misaligned with one or both parties' interests:

1. **Keeping structures as is.** Early in the process, management often pointed to existing operational processes as sufficient, suggesting that new layers or structures for involving workers in AI decisions were unnecessary. However, subgroup members acknowledged that the current structures did not reliably surface workers' concerns early, and that many important conversations occurred before any formal communication about new technologies reached workers or their representatives.
2. **Labor joining existing technology committees.** Labor participants raised this strategy for ensuring earlier involvement. Management countered that many decisions involve financial or confidential topics or require specialized technical knowledge. Both sides recognized that committee access alone would not solve the timing gap.
3. **Putting AI consultations in an existing national task force.** Some participants asked whether AI could simply fall under the Affordability Task Force (ATF) given the ATF's success. Subgroup members rejected this option. They noted that the ATF's cadence and scope would likely crowd out technology issues, and that AI's pace and complexity required a dedicated forum with its own workflow, rather than being absorbed into an already heavy agenda. Our research team had an additional concern that framing AI and technology decisions as primarily an affordability topic might encourage pursuit of technologies that automate tasks and displace more workers, rather than other AI technologies that aim to augment workers' expertise and may displace fewer workers. That concern was not explicitly voiced by subgroup members in our presence, and so we assume it was not part of the bargaining strategy.
4. **Using regional Labor Management Partnership (LMP) Councils.** Regional LMP Councils were considered as potential "middle layer" venues that could drive coordination between the many ideas surfaced by UBTs and the guidance from the Task Force regarding large, multi-region AI projects. However, participants raised concerns about the bandwidth of LMP Councils and regional variation in their operation. LMP Council structures were seen as too variable and overloaded to take on this role consistently. This left the subgroup without consensus on a clear "middle layer."

Negotiations Over Economic Issues

Negotiations by the Common Issues Committee on economic issues began in parallel with the deliberations of the AI subgroup in June 2025 and then continued after the Tentative Agreements on AI were approved by the CIC on August 22, 2025. While our research team did not participate in these economic negotiations, we summarize them in Figure 7 based on public accounts of the process and outcomes. It is clear from the timeline shown in Figure 7 that the negotiations over economic issues were protracted and difficult. Two work stoppages occurred. In negotiations in October, KP presented a wage offer of 21.5% over the course of four years. This offer was ultimately accepted by the unions in March 2026.

In the past, negotiations on economic issues took place both at the national level between KP and the Alliance and with the different unions representing workers in specific locations. In February 2026, however, KP announced it would no longer participate in national-level negotiations and would instead negotiate separate agreements with each local union. Agreements were reached with all of the local unions in March 2026 and then were subsequently ratified by their members. The basic wage agreement called for an increase of 21.5% over the four years of the agreement. The agreements on AI summarized above and listed in the Appendix were incorporated as addendums into each of the local agreements.

The events that led to the breakdown in national negotiations pose significant questions about the future of the labor-management partnership. The good news is the parties have faced and successfully worked through other similar moments of high stress to the partnership over the course of their twenty-nine-year history. In this instance, they signed a Memorandum of Understanding (MOU) stating they would meet within 90 days following ratification to discuss how to “reconstitute” their Partnership.

“Indeed, given the frayed relationships that often carry over from such difficult contract negotiation experiences, working together to implement the provisions of the AI agreement might serve as a means of reminding all stakeholders of the value of partnership.”

“

Figure 7. Negotiations on Economic Issues

Time Period	Activities
June-August 2025	CIC negotiates on economic issues at the national level; local union negotiations begin.
September 2025	National and local negotiations continue to September 30, the date of the contract expiration. No agreement is reached.
October 2025	National and local negotiations continue. The Alliance Unions engage in a five-day strike.
November-December 2025	National and local negotiations continue. No agreement is reached at the national level; there is limited progress in reaching local agreements.
January-February 2026	UNAC/UHCP, a union in the Alliance representing nurses and other clinicians in Southern California and Hawaii, engages in a four-week strike. KP announces an end to national negotiations and states it will continue to negotiate toward reaching agreements with local unions.
February 23, 2026	UNAC/UHCP announces an end to the strike.
February 24, 2026	KP announces that UNAC/UHCP unions accept wage offer of 21.5% over four years.
February-March 2026	Local negotiations continue. Tentative agreements were reached with all local unions in March 2026 and then were subsequently ratified by their members.
March 2026	Parties agree to include the AI agreement approved by the CIC as an addendum to each local agreement.
March 2026	Parties agree to meet within 90 days following ratification to discuss how to reconstitute their partnership.

The parties should be commended for following through and honoring the tentative agreements they reached on AI before the tumultuous events that occurred in bargaining over wages and other issues. It is a tribute to their shared commitment to addressing AI in partnership as called for in the AI agreement. Indeed, given the frayed relationships that often carry over from such difficult contract negotiation experiences, working together to implement the provisions of the AI agreement might serve as a means of reminding all stakeholders of the value of partnership.

Next Steps

Reaching a landmark agreement on AI was just the first step toward bringing workers and the union into AI and making KP an industry leader in the use of AI. In this section, we offer some thoughts on steps that need to be taken to ensure this agreement realizes its full potential in practice.

Including the AI agreements in each of the local union contracts offers opportunities for fostering local and regional experiments. This would be consistent with the “bottom-up” UBT initiatives contained in the AI agreements.

Preparing for the “Bottom-Up” Processes

The decision to use the UBTs as a source for generating ideas from frontline workers for ways AI can improve operations and their jobs has great potential. It does add another responsibility to the UBTs. Care will be needed to ensure the UBTs have the time and capacity to take on this added responsibility. Additionally, as the agreement anticipates, there will need to be resources and updated systems (e.g., an improved UBT Tracker system) to allow for cross-fertilization and learning across the many UBTs and new projects that utilize AI systems and other new technologies.

Members of UBTs would benefit from an initial briefing and ongoing support for generating and processing ideas for using AI. Some basic training on AI should be provided to all UBTs. Given that there are over 3,500 active UBTs at KP, it would be helpful to develop online resources that provide basic information on AI and illustrate the process for generating and submitting ideas for using AI. While online training can provide some orientation, recent research on AI implementation in other healthcare contexts suggests that participatory training and guided exploration is important to help frontline workers see the potential value of different AI systems and the possibility of generative AI for their particular roles.¹² One approach might be to encourage and empower the UBTs to identify problems and opportunities for using AI to address any or all of the points on the Kaiser Permanente Value Compass shown in Figure 8, because the UBTs use the Value Compass to guide their existing work.

Figure 8.
The Kaiser Permanente Value Compass



To support this work, some UBT coaches, who are trained staff who support UBTs using continuous improvement strategies, would benefit from further training to understand existing and possible AI and technology solutions, so they can more effectively help UBTs brainstorm about possible ideas. Additionally, it may be useful to have a subset of IT professionals who specialize in hearing frontline concerns, exploring how those line up with existing AI and technology solutions, and identifying new directions for development work that could address UBT ideas for improving operations, patient care, and employee experience.¹³ In effect, there likely needs to be an investment in “matchmaking” capabilities, so that those leading UBTs develop a better sense of feasible technology solutions and those developing, deploying, and investing in technology solutions are regularly hearing from frontline experts via UBTs.

The subgroup also identified that an improved UBT Tracker could help management, labor leaders, and other UBTs see what had been tried and amplify promising use cases quickly. If the UBT Tracker provides a shared and broadly accessible library of use cases—including fully implemented changes and those that were attempted and stalled—it can be analyzed to help match a given UBT’s concerns or ideas to particular technologies or workflow changes. The UBT Tracker can also help identify local employee leaders who are excited about AI and other technology innovations, as well as cases to highlight in communications campaigns.

One of the lessons learned from introducing automated medical records technologies at KP is that “peer advisors”—union members trained in the use of the new technologies—were valued and trusted sources of assistance to other workers who were experiencing challenges in adapting to those technologies.¹⁴ This suggests that KP and the Alliance might want to train a cadre of AI peers to be trusted resources whom employees can call on as AI tools become incorporated into their jobs. Building this capacity to facilitate the everyday use of AI will be critical to both overcoming the fear or lack of confidence some workers may experience about using AI and helping others overcome their resistance to changes associated with using AI. And, as they do this critical work, the AI peer advisors will be a valuable source of ideas for dealing with problems experienced in the use of AI and of further ideas for adapting AI and work processes to enhance patient care and job quality.¹⁵

During bargaining, KP and the Alliance chose to leave it to the national task force to decide what intermediate group(s) would process suggestions coming from the UBTs or other bottom-up sources and how to learn from all this distributed innovation. This is a key question. Each region will need to put some group in place (or assign this task to an existing group) quickly, and regions would be wise to pay attention to how other regions handle this coordination challenge.¹⁶ Evidence from both other studies of AI¹⁷ as well as decades of experience with other forms of bottom-up idea generation¹⁸ suggest these intermediate bodies are critical for sustaining workforce interest and gaining a steady and well-informed stream of bottom-up ideas.¹⁹

Beyond the UBTs, training and educational materials should be developed and made available to all KP employees to ensure they know the basics about AI, to encourage them to experiment with the use of AI tools, and to encourage them to discuss with their peers their ideas for using AI in the workplace. KP and the Alliance likely need to discuss and then communicate appropriate guidelines for worker experimentation with AI tools to ensure patient privacy and address any other security concerns. This is likely an urgent task since “shadow use” of AI tools (where employees use readily available tools outside of their employer’s systems) is increasingly common and involves clear privacy risks in healthcare settings. The task force is expected to provide guidance on training and communications about privacy and governance, but regional and local guidance is likely needed even before that occurs.

Launching the National AI Task Force

The national AI Task Force called for in the AI agreement will play a critical role in guiding and overseeing all aspects of this agreement. Two specific points regarding the membership of the Task Force are worthy of discussion. First, the management members of the Task Force should include the key executives who decide what technologies to invest in and to either purchase from vendors or develop in-house. This would clearly include executives in IT and perhaps others who work closely with and/or approve significant IT investments. While KP HR or labor relations leaders should also be represented on the Task Force, membership should not be limited to those typically in charge of managing Labor Management Partnership affairs. Second, on the union side, membership should likewise include high-level Alliance leaders and members with deep expertise in AI, either expertise they already have or expertise they are ready to develop now through additional training.

A sensitive aspect of the membership of the Task Force discussed in the AI subgroup was whether and, if so, how to involve one or more leaders of the Coalition, the other major union federation in the KP Labor Management Partnership. The parties did not make a firm decision on this issue but recognized that many of the AI investments would apply to groups that include both Alliance and Coalition members. Working out the specifics of how to involve Coalition representatives in the Task Force and, for that matter, at all levels of the AI processes put in place, is critical work to be done.

One opportunity left to the Task Force is the potential value to be created and/or savings to be achieved in engaging with AI developers and vendors about the problems and opportunities that KP managers and employees identify; that opportunity to pull in developers and vendors was discussed in the subgroup but not explicitly included in the agreements. The presentation about the AFT-OpenAI initiative noted above suggested that AI developers are “hungry” for this type of engagement with potential customers. Even when buying an existing AI tool with an existing model, worker input could inform the design of the interface or how the frontline worker interacts with the tool. For example, it is feasible to build a custom interface for a decision-support tool that prompts the worker to describe their thinking or creates other helpful “friction” in the AI-mediated process, in hopes of reducing overreliance and maintaining skills.²⁰

Engaging the big development companies and the healthcare AI vendors would signal that KP and its labor partners are committed to developing and designing AI tools that address critical problems and opportunities via a “pull” method (i.e., in response to frontline concerns and ideas) rather than only pursuing the tools these outside vendors propose without worker or management input. Left to their own devices, technology developers who interact primarily with executives (who may be focused on cutting costs, including labor costs) are likely to be oriented to automating tasks rather than using AI to augment workers’ performance.²¹ But KP’s initiative could have a significant impact on the development processes used in AI in healthcare, and perhaps beyond.

AI vendors could be further motivated to work with KP in new ways when they see the potential market opportunities (at KP and other healthcare providers) for applications to address problems/opportunities that are identified by the UBTs from the “bottom up.”²² Policy proposals for encouraging “pro-worker AI” have noted that the public sector has great leverage in shaping healthcare technology, as a powerful payer.²³ A complementary and perhaps reinforcing source of leverage could come from the new model that KP and the Alliance pursue through the task force and bottom-up innovation. The big technology companies and the cutting-edge small developers may be pushed to worker-informed developments and pro-worker AI solutions by public-sector requirements and pulled in that direction by innovative, market-defining customers like KP.

A critical duty assigned to the Task Force is to ensure that the effects of this agreement are tracked and assessed to support learning and improvement in timely ways and to engage outside experts to support continuous learning and discussion of AI. Thus, we hope that this case study is just our first installment in an ongoing process of research and shared learning and discussion with KP and its labor partners on the work ahead.



Broader Implications of This Case

Because there are few labor-management partnerships in the U.S., some might argue that KP and the Alliance represent a very special case that has few implications or little relevance for other unionized settings in the U.S. and no relevance for the country's many nonunion workplaces.

We agree that this setting is unusual in the U.S., but we think it has significant implications for other settings. In this last section, we outline six lessons this case might suggest to those who share the view that bringing workers' voices into AI decisions is an important step for assuring that AI is used responsibly and that the benefits of these new technologies are shared equitably.

1. First, employers and labor unions shouldn't dismiss the power of interest-based, problem-solving approaches to negotiations.

KP has had a twenty-nine-year history of labor-management partnership and the parties used an interest-based problem-solving subgroup process to address AI issues within a larger, quite contentious negotiation of a new collective bargaining contract. The parties could not have achieved the results summarized in this report without the problem solving built into their negotiations' processes. We witnessed the best features of sustained dialogue that focused on interests, options, and, where possible, consensus decision-making throughout the more than 150 hours of subgroup deliberations. The parties engaged in genuine dialogue and learning, gained a deep understanding of each other's perspectives, and demonstrated mutual respect. This could not have been achieved through traditional position-based negotiations.

To be sure, there were elements of hard positional bargaining in the subgroup discussions and instances of frustration within the union and the management teams. But the parties worked through these difficult moments to reach an agreement that largely reflected their shared interests and their shared commitment to the mandate stated in the subgroup's charter.

At the same time, our research team observed some factors that limited the problem-solving potential of the discussions. There were relatively few management representatives at the table who make decisions about what technologies KP now chooses to design, purchase, or roll out. This led some of the union participants to wonder who, or what higher-level body, makes such decisions. In addition, some management representatives were wary of adding any more joint committees to the numerous labor-management forums already in place. Some union representatives expressed frustrations in private conversations that some of their management counterparts simply wanted to hold the line on new committees but were reluctant to express this as one of their explicit interests. In turn, some of the management representatives expressed in private conversations with the research team that they saw little value being generated in some of the existing labor-management joint efforts and therefore were not convinced that it would be different on AI issues. These unexpressed views served as an undercurrent to the on-the-record conversations. This slowed the progress toward reaching consensus on how to move forward together.

Undoubtedly, some of these management and labor concerns will surface again as the parties implement the terms of their agreement. It will be important for all involved to surface their doubts and concerns and to work to ensure that the structures and processes called for in the agreement are indeed productive and that appropriate steps are taken to follow up on recommendations and/or decisions that are generated in these labor-management forums.

To be sure, interest-based, problem-solving approaches to negotiations are not a substitute for hard bargaining by either labor or management when hard bargaining is needed to advocate for an interest resisted by the other party. Interest-based bargaining can help open up new ground, as in this case, on nonmandatory subjects of bargaining. But joint problem solving requires that decision-makers are “in the room” and actively engaged in making agreements that they then are committed to supporting and implementing.

2. There is real power in labor solidarity.

A broad mix of unions came together in the Alliance, fashioned a shared vision on AI, and supported each other in the subgroup discussions and in their private caucuses. None of the individual unions would have been able to achieve this on their own. The labor movement has bifurcated into separate national and local union interests all too often. In this case, unions representing many professions coordinated around their shared vision and brought together otherwise fragmented knowledge and experiences. Solidarity is needed for breakthroughs like this and, as we will point out below, will be needed even more for breakthroughs in other sectors and organizations.

3. The broader labor movement can learn from the distinctive approach pursued by KP and the Alliance.

As noted below, other unions and the AFL-CIO have pursued critical AI protections in recent bargaining, but this labor-management partnership agreement provides an example of broader, more consistent involvement of both labor leaders and frontline workers during the term of the labor agreement. It provides an unusual example of more grounded, regular collective voice on AI design, development, and implementation than we have seen in other agreements.

4. Conversely, KP and the Alliance would be wise to monitor what other major labor organizations and business leaders, both in the U.S. and around the globe, are doing to address AI issues.

The KP-Alliance agreements emphasize continued collaboration and ongoing consultation and innovation infrastructure—provisions we see as innovative and exciting—but these agreements give less attention to guardrails to protect employees; this perhaps reflects the fact that local agreements are more focused on those topics. Nonetheless, the partners will want to watch and learn from other negotiated agreements.

Two notable examples are the company-wide agreements between UNITE HERE and major casino and hotel operators and a similar UNITE HERE agreement with Aramark. The company-level agreement between Aramark and UNITE HERE provides a framework for guiding local-level negotiations over uses of new technologies in the many convention centers, ballparks, and other unionized entertainment sites the company manages. Additionally, two highly publicized sector-wide agreements governing uses of AI were reached with the Alliance of Motion Picture and Television Producers in 2023, after two long strikes, by SAG-AFTRA and the Writers Guild of America. In addition, the state of Pennsylvania and SEIU agreed to a pilot project and follow-up evaluation of how state employees might use AI to improve their work and the services they provide; this approach is perhaps most similar to the KP-Alliance agreement in that it provides a framework to dig into changes in work processes. In addition, the AFL-CIO and several national unions such as the Communications Workers of America and the American Federation of Teachers have established joint programs with several of the largest AI vendors, including Microsoft and OpenAI, to educate members about AI and to share ideas for developing AI tools that improve their jobs and the services members deliver.

Judging from the attention the labor movement is now giving to these issues, this list is likely to continue to grow in the years ahead. For example, in March 2026, the AFL-CIO and the AFL-CIO Technology Institute hosted a “Workers First AI Summit” that brought together union leaders, policy experts, academics, and foundation leaders to discuss current developments and future strategies for addressing worker concerns about AI. Participants paid special attention to the need for stronger guardrails governing issues of data privacy, surveillance and monitoring of workers, and the use of algorithmic tools in human resource management decisions. They called for state and national policies to regulate these issues as well as for unions and employers to take up these issues as top priorities in collective bargaining. The broad range of unions participating in the Summit suggests that this is indeed an issue rising to the top of agendas in bargaining and policy advocacy.

KP, the Alliance, and the U.S. labor movement can also learn from their counterparts abroad, particularly from the policies and institutions supporting worker voice and representation in European Union (EU) countries. Indeed, the EU and its member nations are generally viewed as being far ahead of the U.S in addressing AI. The EU has enacted two relevant sets of regulations: the General Data Protection Regulation (GDPR) in 2016 and the AI Act in 2024.²⁴ Both set limits on the use of personal data, and the AI Act specifies requirements for transparency in the uses of AI and for managing the risks of different forms of AI. Much of the enforcement of these rules, however, is left to individual countries.

European labor unions also appear to be ahead of their U.S. counterparts in gaining a voice for workers on AI. Within countries such as Norway, Finland, Germany and Spain, unions have negotiated national and/or sector-wide agreements governing AI and provide workers the right to form “works councils” (worker representation bodies that must be consulted prior to the implementation of new technologies).²⁵ These and other efforts to bring workers’ voices into AI policies and practices received another boost in Pope Leo XIV’s recent encyclical on the need for a “human-centered” approach to AI development and use.²⁶ Thus, there are developments on the international stage that are worth monitoring as parties in the U.S. explore how to strengthen AI policies and practices.

5. There are benefits for employers to having worker engagement and voice structures in place.

Here we turn to the question of what both KP executives and other management leaders might take away from this case. One core insight is that organizations aspiring to lead in AI and to achieve “AI transformation” should engage all parties. At a minimum, meaningful engagement will be needed to overcome the fear and skepticism of frontline workers (for example, as expressed by rank-and-file union members at the pre-bargaining workshops). Absent structures and processes for getting worker knowledge and ideas into the development and design phases illustrated in Figure 2, implementations of top-down “solutions” would very likely meet resistance, continue to require expensive redesign, and delay the time involved in achieving significant returns on technology investment decisions. KP is able to leverage the labor-management partnership infrastructure—frayed though it may be—and draw on the structures, processes, and mindsets from previous partnership efforts to improve the organization and address critical challenges.

Other organizations, particularly those whose workers are not currently represented by a union, will need to invest in this kind of employee participation and worker voice infrastructure in order to engage frontline workers and avoid the “last-mile” implementation challenges common in AI efforts.²⁷ While many organizations discuss the value of worker input in their AI efforts, those contributions can be framed as the strategic sourcing of critical frontline expertise and tacit knowledge that will allow AI systems to work—not as an opportunity for those workers to shape the design or implementation of new AI tools or to influence how such tools are

integrated into their workflows. Organizations too often see workers as “end users” who need to be consulted as new technologies are introduced to see if they work. But the critical decisions such as the definition of what AI is designed to do, the actual design and adaptation of the AI tools, and the workflow and process changes are made by others, before the workforce is consulted and without an orientation to AI augmenting human work. Some organizations also are promoting and supporting bottom-up idea generation on AI use cases—but they may not have the psychological safety to actually elicit those ideas or the management commitment to learn from frontline employees’ insights.

Could these firms learn from the KP-Alliance agreements? Absolutely. These processes could be adapted to a variety of organizations. However, the question would be: will frontline input be meaningful in situations without union representation or some means of exercising collective voice? In these settings, what conditions or strategies overcome traditional management resistance to giving workers a voice in technology decisions, which are typically viewed as managerial prerogatives?

6. It will be important for KP and the Alliance to monitor and learn from how others are addressing these issues in this fast-changing field, today and in the future.

There is no “one best way” for engaging these issues, so the aim should be to continue documenting these different approaches and strategies and to share the lessons learned widely with employers, workers, unions, AI developers, and AI researchers.

Fortunately, research on these issues is expanding at a very rapid rate. Our own research group at MIT, the Institute for Work and Employment Research, plans to continue studying and working in partnership with companies and unions (in healthcare and other industries) to document ways to use AI tools to improve the quality of goods and services and the quality of jobs. Other MIT colleagues, including researchers at the Stone Center on Inequality and the Future of Work, have been explicit in studying ways to advance “pro-worker AI.”²⁸ The Human-Computer Interaction Institute at Carnegie Mellon University works directly with unions and companies to design and implement AI tools that augment workers’ jobs.²⁹ The UC Berkeley Labor Center advises state governments and labor organizations on legislation regarding guardrails for the use of AI³⁰ and has published an extensive list of collective bargaining clauses addressing technological change.³¹ The ILR School at Cornell is the leading source of information on what U.S. unions and companies can learn from how European unions, works councils, employers, and governments are addressing AI issues, often in ways that are more advanced than is the case in the U.S.³² Colleagues at Penn State’s AI Learning Collaborative host webinars and advise labor leaders in healthcare and other industries on how they can work on AI-related issues.³³ These are just samples of the growing body of research underway that can help inform KP and the Alliance as they address the challenges and opportunities that lie ahead in implementing this agreement.



Moving Forward

In the economic negotiations that took place leading up to the 2026 contract, KP and the Alliance and affiliated unions experienced the most bitter and difficult breakdowns in bargaining over wages and other economic conditions in the twenty-nine-year history of the Labor Management Partnership. As a result, some might question whether partnership is still a viable and valuable strategy for either KP or the unions that represent the workforce. The good news is that the parties recognized the need to address this question and have signed a Memorandum of Understanding to begin a dialogue over how to “reconstitute” the Labor Management Partnership. We view this as a necessary and critical first step toward repairing and rebuilding the trust that will be needed to invent and implement the features of a renewed Partnership—hopefully one that builds on the breakthrough AI agreement documented in this report. Indeed, implementing the AI agreement from the “bottom up” via initial pilot projects or experiments with oversight by a newly created national AI Task Force may serve as a path forward to achieving the shared benefits that will be essential to rebuilding support for partnership among its skeptics. We believe the potential to do so exists, and we look forward to continuing to document and support these efforts by working with all who have a stake in KP’s success, from frontline workers to top executives and physicians to local union representatives and leaders of the Alliance, the Coalition, and other KP unions.

Appendix

Agreement on Artificial Intelligence Approved by the Common Interest Committee, August 22, 2025

AI and Technology

Kaiser Permanente and the Alliance of Health Care Unions recognize that AI and evolving technologies are transforming healthcare. We are committed to the use of AI and technology to improve health outcomes and the patient, member, and employee experience. Working in partnership, we strive to ensure that new technology serves our patients, members, employees and the organization. KP and the Alliance are committed to timely communication to the workforce about new technology that impacts employees' roles and workflows.

The parties agree to the following:

Relational care:

Compassion, communication, human touch and relational care are essential to delivering high-quality healthcare. AI and technology are tools to improve patient-centered care, outcomes, and employee work experience. AI and technology shall harness the expertise of frontline workers, and complement clinical judgment and professional practice with human evaluation and the ability to intervene, as appropriate.

Employee voice:

The employee voice is essential to the successful adoption of new technology. Through the agreed-upon Alliance-KP Taskforce on Technology and Artificial Intelligence and UBT structures, employee input will be incorporated into the AI and technology lifecycle, which may include but is not limited to: problem identification and solution design to implementation, education, training, evaluation, and communication.

Ethical use of AI and Technology:

To ensure the ethical, safe and principled deployment of technology, it is essential to avoid discrimination, bias, or inequity, and to uphold transparency and accountability. The implementation of AI and technology shall reflect partnership values.

The employer will inform the union, per local practice, of any new AI technology that will be implemented for monitoring productivity and/or performance. Disputes under this paragraph shall be resolved exclusively pursuant to the applicable local collective bargaining agreement(s), and shall not be subject to any dispute mechanisms under this National Agreement.

KP and the Alliance of Healthcare Unions have previously agreed to guidance that addresses the impacts of AI and Technology on employees, e.g., Union Security (Section 1.K.5 of this agreement), the Employment Income Security Agreement, and Workforce Planning and Development (Section 1.D of this agreement).

The parties have agreed to create and utilize a taskforce on technology and artificial intelligence and the UBT structure to capture the employee voice and expertise. We reiterate our commitment to the use of AI and technology to improve health outcomes and the patient, member, and employee experience.

AI & Technology Subgroup Recommendation (Approved by CIC 8-6-2025)

Alliance-KP Taskforce on Technology and Artificial Intelligence

I. Timeline

The parties will establish the Alliance-KP Taskforce on Technology and Artificial Intelligence within 120 days of ratification.

II. Task Force Objectives

- **Share** information about planned innovation, technology, and artificial intelligence rollouts
- **Learn** together about how technology and artificial intelligence are impacting KP and healthcare more broadly
- **Guide** communication, education, and training related to technology and artificial intelligence

III. Task Force Elements

A. Scope

- Share and discuss National initiatives related to technology and artificial intelligence (defined as initiatives impacting two or more Markets)
 - The task force may agree to review other significant large-scale initiatives brought forward by either party
- Partner to support and discuss rollouts and trouble-shoot implementation on National initiatives
- Discuss current technology utilization and prospective technology adoptions
- Support a communications strategy to inform Labor and Management around the value of partnering on technology
- Identify education and training needs for the workforce
- Hold strategic discussion about potential future innovation and technology adoption
- Serve as a working body, not think tank or discussion forum, that works with (does not replace) existing forums/groups
- Discuss and bring in resources and subject matter experts, including:
 - from academia, for training, feedback and evaluation
 - resources in support of this technology and AI work
 - experts in regulatory and legislative affairs
- Provide a forum for early engagement to make recommendations for effective development, implementation, communications, and/or training
- Discuss topics of importance or concern to either Labor or Management related to technology and Artificial Intelligence, which may include hiring, performance evaluation and corrective action.
- Discussion at the Task Force does not supplant local engagement (impact notification and bargaining)
- Develop recommendations to educate on the importance of partnership in AI & Technology throughout the enterprise
- Make recommendations for effective market level structures based on escalations from a market

B. Membership

- 10 members (5 labor and 5 management), with appropriate participation of additional attendees/invitees.
- Intent is for core membership to include senior leadership from labor and operations, executive leadership with experience in technology (including but not limited to artificial intelligence), innovation, operations, learning

IV. Deliverables

A. Identify and Recommend Training Needs

- Providing input on training needs, including but not limited to:
 - gaps in training needs
 - skills as a whole
 - training for particular pieces of technology associated with a technology rollout
- The task force would not create the training plans but could provide guidance for training needs to the people/departments that are creating the training plans

B. Be a Connective Hub for Technology and AI Topics

- Make connections between entities related to innovation, technology, and AI readiness (i.e., Ben Hudnall, Enterprise Learning, LMP Learning, etc.) in order to identify gaps, reduce overlap, identify and leverage existing programs

C. Workforce Communications

- Informing, coordinating, and/or advising on a communication strategy regarding technology and AI
- Provide content and guidance to LMP Communications and KP National Communications to distribute information to the workforce through various communication channels

V. Expected Outcomes

- Increased communication and collaboration related to innovation, technology, and artificial intelligence
- Successful, accelerated, and sustainable implementation of technology and artificial intelligence solutions
- Promote workforce engagement in development, adoption, and rollout of AI and Technology for improved patient care, efficient operations, and improved work experience.
- Engaged and future-ready workforce adept at innovation and technology adoption.
- Recommendations related to development, implementation, communications, training, and/or improved work experience.

VI. Evaluation and Continuous Improvement

- Taskforce will develop methods for ongoing assessment of proposal effectiveness
- Taskforce will develop plans for adapting and refining the approach based on results

Enhancing the Unit-Based Team Structure to Amplify Frontline Worker Voice and Drive Innovation

Leveraging Existing Teams for Innovation and Service Excellence

I. Proposal Objectives

- **Empower** frontline workers to contribute directly to innovative solutioning
- **Enhance** support mechanisms for innovation within UBTs
- **Establish** a clear pathway for innovative ideas to scale across business units and bi-directional communication.

II. Strategic Focus Areas and Implementation Plan

A. Employee Education and Communications

- Ensure inclusion of national innovation training into core UBT training series (Path to Performance) and implement initial and ongoing training initiatives.
- The LMP Council KP and Alliance Members will ensure there are methods and best practices to highlight innovation success stories and lessons learned.
- The LMP Council KP and Alliance Members will evaluate communication channels to ensure timely, transparent information flow in regards to innovation between UBTs, co-leads and sponsors, and local and/or regional LMP councils.

B. Designation of Market Responsibilities

- The local and/or regional LMP Council KP and Alliance Members will designate responsibility to implement these provisions of the KP-Alliance National Agreement for the following no later than April 1, 2026.
 - o Provide tools and templates to guide role execution in supporting frontline driven innovation
 - o Develop mentorship and support systems to build capacity and confidence in the roles supporting innovation initiatives
 - o Responsibility to review and report innovation UBT projects and ideas
 - o Dissemination of National and Regional updates and initiatives
 - o Coordinate success story strategy

C. UBT Tracker Improvements

Business Requirement Considerations:

- Ability to identify UBT tech innovation opportunities
 - o (e.g. capture ideas and not just sponsored projects in the tracker)
- Ability to pull insights on tech innovation opportunities, trends, progress within each region and across the program
- Ability to see real-time data, visualizations (e.g., dashboard)

III. Expected Outcomes

- Increased engagement and empowerment of frontline staff
- Accelerated identification and scaling of innovative solutions
- Improved service delivery and operational efficiencies

Notes

¹ National Academies of Sciences, Engineering, and Medicine, *Artificial Intelligence and the Future of Work* (Washington, DC: The National Academies Press, 2025), <https://doi.org/10.17226/27644>.

² Thomas A. Kochan, Adrienne E. Eaton, Robert B. McKersie, and Paul S. Adler, *Healing Together: The Labor-Management Partnership at Kaiser Permanente*. (Ithaca, NY and London: Cornell University Press, 2009); and Thomas A. Kochan, *The Kaiser Permanente Labor Management Partnership: 2009–2013* (Cambridge, MA: MIT Institute for Work and Employment Research, 2013).

³ Robert B. McKersie, Teresa Sharpe, Thomas A. Kochan, Adrienne E. Eaton, George Strauss, and Marty Morgenstern, "Bargaining Theory Meets Interest-Based Negotiations: A Case Study," *Industrial Relations* 47, no. 1 (January 2008): 66–96, <https://doi.org/10.1111/j.1468-232X.2008.00504.x>.

⁴ Adam Seth Litwin, "Technological Change at Work: The Impact of Employee Involvement on the Effectiveness of Health Information Technology," *ILR Review* 64, no. 5 (October 2011): 863–888, <https://doi.org/10.1177/001979391106400502>; and Kochan et al., *Healing Together*.

⁵ Federal Mediation and Conciliation Service, "Interest-Based Bargaining: Joint Problem-Solving for Mutual Gain," n.d., <https://www.fmcs.gov/wp-content/uploads/2026/04/Interest-Based-Bargaining-2026.pdf>.

⁶ Kochan et al., *Healing Together*.

⁷ Ibid.

⁸ Albert Rees, "Industrial Conflict and Business Fluctuations," *Journal of Political Economy* 60, no. 5 (October 1952): 371–82, <https://doi.org/10.1086/257272>.

⁹ Katherine C. Kellogg, Mark Sendak, and Suresh Balu, "AI on the Front Lines," *MIT Sloan Management Review* 63, no. 4 (Summer 2022): 44–50, <https://sloanreview.mit.edu/article/ai-on-the-front-lines>.

¹⁰ Ben Armstrong and Benjamin Berkowitz, "Scaling Automation: Two Proven Paths to Success," *MIT Sloan Management Review* 65, no. 2 (Spring 2024), <https://sloanreview.mit.edu/article/scaling-automation-two-proven-paths-to-success>.

¹¹ Thomas A. Kochan, Ben Armstrong, Julie Shah, Emilio J. Castilla, Ben Likis, and Martha E. Mangelsdorf, "Bringing Worker Voice into Generative AI," *An MIT Exploration of Generative AI*, March 2024, <https://doi.org/10.21428/e4baedd9.0d255ab6>.

¹² Kellogg, Sendak, and Balu, "AI on the Front Lines."

¹³ Armstrong and Berkowitz, "Scaling Automation."

¹⁴ Litwin, "Technological Change at Work."

¹⁵ Wanda J. Orlikowski, "Sociomaterial Practices: Exploring Technology at Work," *Organization Studies* 28, no. 9 (September 2007): 1435–1448, <https://doi.org/10.1177/0170840607081138>.

¹⁶ Kellogg, Sendak, and Balu, “AI on the Front Lines”; and Armstrong and Berkowitz, “Scaling Automation.”

¹⁷ Kellogg, Sendak, and Balu, “AI on the Front Lines”; and Arvind Karunakaran, “AI in the Workplace: Rethinking Skill Development,” Webinar, Stanford Online, September 5, 2024, <https://www.youtube.com/watch?v=U4QWzTPg1gl>.

¹⁸ Frederick G. Lesieur, ed. *The Scanlon Plan: A Frontier in Labor-Management Cooperation* (New York: John Wiley & Sons, 1958).

¹⁹ Batia M. Wiesenfeld and Katherine C. Kellogg, “Generative Organizational Learning: Affordances for New Modes of Knowledge Search, Creation, Transfer, and Forgetting with LLMs,” *Strategic Organization* 24, no. 2 (May 2026): 256–284, <https://doi.org/10.1177/14761270251408580>.

²⁰ Zana Bućinca, Maja Barbara Malaya, and Krzysztof Z. Gajos, “To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-Assisted Decision-Making,” *Proceedings of the ACM on Human-Computer Interaction* 5, Issue CSCW1, Article no. 188 (April 2021): 1–21, <https://dl.acm.org/doi/10.1145/3449287>; and Zana Bućinca, Siddharth Swaroop, Amanda E. Paluch, Finale Doshi-Velez, and Krzysztof Z. Gajos, “Contrastive Explanations That Anticipate Human Misconceptions Can Improve Human Decision-Making Skills,” *CHI ’25: Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, Article no. 1024 (April 2025): 1–25, <https://dl.acm.org/doi/10.1145/3706598.3713229>.

²¹ Kochan et al., “Bringing Worker Voice into Generative AI”; and Daron Acemoglu, David Autor, and Simon Johnson, “Building Pro-Worker Artificial Intelligence,” National Bureau of Economic Research Working Paper 34854 (February 2026), <https://doi.org/10.3386/w34854>.

²² Jenna E. Myers, “Triadic Technology Configuration: A Relational Perspective on Technologists’ Role in Shaping Cloud-Based Technologies,” *ILR Review* 77, no. 3 (May 2024): 307–335, <https://doi.org/10.1177/00197939241232992>.

²³ Acemoglu, Autor, and Johnson, “Building Pro-Worker Artificial Intelligence.”

²⁴ Sally Wright, “Artificial Intelligence and Work: A Review of the European Policy Landscape,” *Journal of Industrial Relations* 67, no. 5 (November 2025): 794–805, <https://doi.org/10.1177/00221856251394780>.

²⁵ Virginia Doellgast, Tobias Kämpf, and Barbara Langes, “Building Worker Voice and Power in AI Decisions: Three Cases in the German ICT Industry,” *Work and Occupations* (2026), <https://doi.org/10.1177/07308884251412886>.

²⁶ Leo XIV, “Magnifica Humanitas,” encyclical letter, May 15, 2026, <https://www.vatican.va/content/leo-xiv/en/encyclicals/documents/20260515-magnifica-humanitas.html>.

²⁷ Yuri Nakao, Simone Stumpf, Subeida Ahmed, Aisha Naseer, and Lorenzo Strappelli, "Toward Involving End-Users in Interactive Human-in-the-Loop AI Fairness," *ACM Transactions on Interactive Intelligent Systems (TiiS)* 12, no. 3, Article no. 18 (July 2022):1-30, <https://doi.org/10.1145/3514258>; Jack Gallifant et al., "A Field Guide to Deploying AI Agents in Clinical Practice," working paper (December 2025), <https://ssrn.com/abstract=5732262>; Kellogg, Sendak, and Balu, "AI on the Front Lines"; and Karim R. Lakhani, Jared Spataro, and Jen Stave, "The 'Last Mile' Problem Slowing AI Transformation," *Harvard Business Review*, March 9, 2026, <https://hbr.org/2026/03/the-last-mile-problem-slowing-ai-transformation>.

²⁸ Acemoglu, Autor, and Johnson, "Building Pro-Worker Artificial Intelligence."

²⁹ Franchesca Spektor et al., "Working Together: Algorithmic Management and Peer Relationships in the Hospitality Industry," in *DIS '25: Proceedings of the 2025 ACM Designing Interactive Systems Conference* (July 2025): 3221–3234, <https://doi.org/10.1145/3715336.3735704>.

³⁰ Annette Bernhardt, Lisa Kresge, and Reem Suleiman, "The Data-Driven Workplace and the Case for Worker Technology Rights," *ILR Review* 76, no. 1 (January 2023): 3–29, <https://doi.org/10.1177/00197939221131558>.

³¹ Lisa Kresge, "Negotiating Tech: An Inventory of U.S. Union Contract Provisions for the Digital Age," (UC Berkeley Labor Center, updated July 2, 2025), <https://laborcenter.berkeley.edu/negotiating-tech>.

³² Virginia Doellgast, Shruti Appalla, Dina Ginzburg, Jeonghun Kim, and Wen Li Thian, *Global Case Studies of Social Dialogue on AI and Algorithmic Management*, International Labour Organization working paper 144 (July 2025), <https://doi.org/10.54394/VOQE4924>.

³³ Pennsylvania State University School of Labor and Employment Relations, "Learning Collaborative on Artificial Intelligence (AI) for Healthcare Unions and Union Leaders," n.d., <https://ler.la.psu.edu/labor-school/learning-collaborative-on-artificial-intelligence-for-unions-and-union-leaders/>.