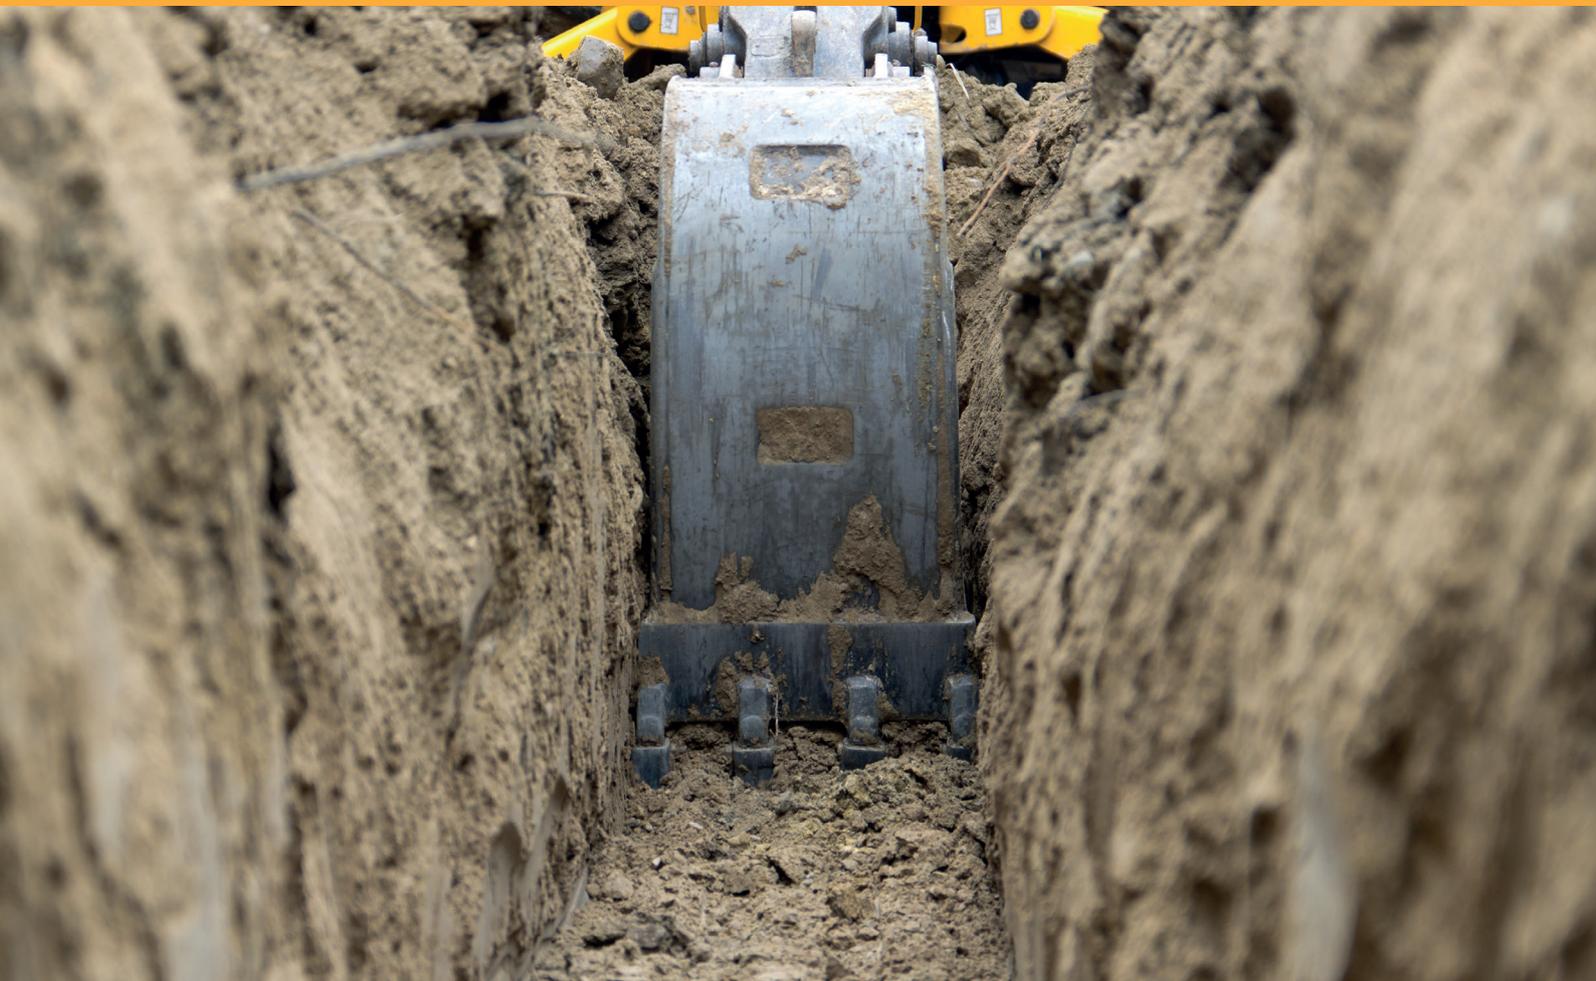


Created for JCB by UKG

Workforce Management in the Age of AI



Reporting Results from the 2018
Global Retail Absence Survey by
The Workforce Institute at UKG

Executive summary

Artificial intelligence¹ (AI) and machine learning² (ML) techniques are disrupting industries and products from thermostats to automobiles. According to Gartner, the global enterprise value derived from AI will total \$1.2 trillion this year, a 70 percent increase from 2017. This figure is expected to increase to \$3.9 trillion by 2022. But what do these techniques mean to the field of workforce management? This document discusses the UKG experience with workforce management AI, particularly in the Workforce Dimensions product. First, it describes the AI best practices that UKG uses, covering the lifecycle from matching a business problem with an AI technology through testing and delivery. Next, this document covers UKG-specific AI advantages, including:

- **Utilising big data with the UKG D5™ architecture** to find patterns hidden across an organisation
- **Exploiting UKG knowledge** of features and metrics that drive workforce management
- **Testing and refining with real data** based on customer partnerships and research

Detailed benefits of using AI in specific workforce management case studies are covered at the end of the document. Examples of how AI can make workforce management easier:

- **Timecard compliance auditing** can be significantly streamlined by using unsupervised machine learning, potentially saving millions of pounds in lawsuits
- Business **volume forecasting** using supervised machine learning can improve the accuracy of forecasts by 20 percent, an increase worth tens of millions of pounds as schedules better reflect demand
- **Shift swap recommendations** can be made using optimisation and preference learning techniques, enhancing employee control of schedules and leading to better engagement

Introduction: AI, machine learning and workforce management

This independent research programme was conducted on behalf of The Workforce Institute at UKG Incorporated by Coleman Parkes Research, an independent U.K.-based research company. Survey data was collected in June and July 2018 from 800 respondents using an online quantitative methodology. Survey participants were sourced from multiple global markets and represent a variety of retail sectors, including grocery, department stores, warehouse, specialty, convenience, and discount. Survey participants included retail managers, store managers, and heads of store operations

What is AI?

Artificial intelligence is not a single monolithic algorithm or even a particularly unified field. It is a collection of subfields that attempt to solve problems traditionally associated with human intelligence. These subfields include computer vision, natural language processing, automated planning and optimisation, and machine learning. In addition, each of these subfields has classes of algorithms matching particular kinds of problems.

¹S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed., Upper Saddle River, NJ: Prentice Hall (2010).

² Christopher M. Bishop, *Pattern Recognition and Machine Learning*, New York: Springer (2006).

For instance, consider machine learning, which generally covers approaches that extract patterns or models from data. Within ML, there are many classes of algorithms, including:

- **Unsupervised machine learning**, where data is analysed to extract features or patterns that best generalise or describe the data. Examples include clustering algorithms and component analysis algorithms.
- **Supervised machine learning**, where each row of data has an associated label (in a classification task) or number (in a regression task) that the model is meant to predict. For instance, a predictor of daily business volumes would have historical volumes to help calibrate, or “train,” it. Algorithms for solving these types of problems include deep neural networks, tree ensembles, support vector machines, Gaussian processes, and many other techniques.
- Other techniques or variants of the above include **semisupervised learning** (where only some of the data has labels), **preference learning** (where the predicted labels are hidden and must be inferred from choices or direct queries), and **reinforcement learning** (where an AI agent makes sequential decisions with the intention of minimising costs).

This rich landscape of solutions is both a wealth of untapped potential and a potential headache for practitioners if their business problem is ill-defined. To build AI business solutions in workforce management or in any other domain, it is imperative to have a well-defined business problem and then match it to the right subfield(s), class, and algorithm.

AI and workforce management business problems

Workforce management covers a wide array of operations: timekeeping, scheduling, volume forecasting, payroll, attrition, hiring, and many other aspects of modern labour logistics. While classical AI approaches have been used for a long time in some parts of this domain (for instance, stochastic search techniques used to generate schedules), the adoption of more modern AI and ML techniques has now begun. Recent surveys show that while much uncertainty remains about their eventual impact,³ AI is anticipated to automate and streamline processes to everyone’s benefit.⁴

To ground this discussion and show the quantifiable benefits of these techniques, this document focuses on three case studies where AI can improve workforce management:

- **Timecard compliance** — With potentially millions of timecard edits and exceptions per year, it can be difficult for organisations to find nuanced compliance violations, such as managers subtly changing break times to avoid meal penalties. Unsupervised learning approaches such as clustering can detect the general patterns and uncover needles in this haystack of audit data.
- **Business volume forecasting** — The first step in any optimised scheduling solution is to predict the amount of business demand on a given day. This type of forecasting is a prototypical use case for supervised machine learning regression, which can make more accurate predictions utilising many more factors and data points than traditional formulas can.

Machine learning vs. deterministic rules:

In some cases, using deterministic or rules-based AI is the right approach. This is especially true in relation to compliance and an organisation’s best practices. At UKG, we understand that there are times when either a rules-based or a machine-learning approach is the best fit, and there are times when a combination of the two is better suited. Reminding employees to take vacations, alerting administrators about employees approaching FMLA leave status, and diagnosing scheduling or attendance issues often require applying algorithms with set rules that don’t require learning and inferring from historical data. The results of deterministic rules are always 100 percent accurate. On the other hand, predicting turnover, absenteeism, or flight risk based on historical data or patterns with an acceptable degree of statistical certainty — almost never at 100 percent — will require machine learning techniques.

³ Erin Winick, “Every Study We Could Find on What Automation Will Do to Jobs, in One Chart,” MIT Technology Review (January 25, 2018), found at <https://www.technologyreview.com/s/610005/every-study-we-could-find-on-what-automation-will-do-to-jobs-in-one-chart/>.

⁴ “Can Artificial Intelligence Make Work Better?” Workforce Institute at UKG (February 20, 2018), found at <https://workforceinstitute.org/artificial->

- **Employee control of scheduling** — Shift “swaps,” where an employee swaps a shift he or she cannot work with another employee’s shift, are common practice across industries. But finding the right shift in a sea of possibilities can be a daunting task. An AI solution based on preference learning that can recommend potential shifts to employees based on their inferred scheduling preferences can make this process easy and efficient.

These areas are not only core aspects of workforce management but also high-value use cases. Finding timecard compliance issues early could save millions of pounds in wage manipulation and tens of millions of pounds in potential lawsuits. ML volume forecasting improvements lead to more accurate schedules, representing millions of pounds in savings for large retailers. And providing an AI-enabled interface where workers can quickly take control of their schedules will lead to less wasted time and higher employee engagement.

UKG AI best practices and differentiators

Delivering a true and reliable business solution takes more than the latest fancy AI algorithms. To ensure that AI and ML techniques can meet their potential in workforce management, UKG has developed a set of AI best practices that apply across its solutions:

- **Match** specific AI solutions to business problems with real-world value. There is no single monolithic “AI algorithm,” and UKG has been successful at pairing AI and business problems such as compliance (unsupervised ML), volume forecasting (supervised ML), and shift swapping (with preference learning and optimisation).
- **Innovate** by combining AI with decades of UKG workforce management knowl-edge. Utilising this knowledge to design features or scalability is often the source of true differentiators for an AI solution.
- **Test** models on real data through customer partnerships. These tests ensure both accuracy and alignment with customer needs.

- **Deliver** understandable and actionable results to users. Delivery includes not only a usable front end but also a scalable back-end architecture.

UKG differentiators

The best practices listed above are good academic guidance, but to take these lessons out of the laboratory and into a real-world product means building out significant infrastructure and processes. UKG offers real-world differentiators in the AI space, including:

- **Utilising big data with the UKG D5 architecture** — The UKG D5 architecture running in the cloud allows Workforce Dimensions to scale up computational resources to meet the needs of modern AI approaches and exploit enterprisewide data. For instance, the AI timecard compliance solution is able to analyse an organisation’s entire data set to find the departments that are truly different.

- **Exploiting UKG knowledge** — Arguably, the most important part of any AI solution is choosing the features and metrics to input, including their representation. This takes knowledge of not just what the AI algorithms do but also the business problem at hand. What are the 20 most important features for predicting business volume two weeks from now? What six metrics best capture timecard compliance risk? In Workforce Dimensions, UKG brings over 40 years of workforce management experience to bear on these crucial questions.
- **Testing and refining with real data** — As mentioned in the best practices above, UKG partners with its customers to test AI solutions on real data before moving to production. This provides critical tests and checks on the design and quantifiable confidence in the real solution. For instance, the volume forecasting solution described below was tested on many different retailers and yielded a median of 20 percent accuracy improvements over existing methods.

Workforce management AI case studies

To demonstrate the power and value of AI solutions for workforce management problems, three case studies are described below. They cover timecard compliance, business volume forecasting, and employee control of schedules. Each case pairs the business problem and AI solutions and then describes UKG-engineered innovations and advantages as well as the testing and delivery methods.

Timekeeping compliance

Uncovering fraud and compliance is a common AI use case across industries. For instance, credit card fraud is often detected with AI algorithms trained on typical customer usage patterns. While one transaction may not be cause for alarm, multiple transactions will usually trip the alarm. Workforce management has a similar issue in timecard compliance. Timecard compliance means ensuring that all punch editing, schedule manipulation, and timekeeping exceptions follow legal and organisational regulations and best practices. While most companies utilise rules to enforce legal constraints, more subtle behaviors can still result in large compliance

violations. For instance, it is common for managers to make small changes to employee punches, such as extending a break from 20 minutes to 30 minutes, to avoid violations. Just as in the credit card example, while one such edit may be rationalised as fixing a simple input error, systematically making similar edits implies a clear violation.

Unfortunately, most large organisations have hundreds of thousands, if not millions, of timecard edits and exceptions each year. Auditing all these data points by hand is impractical. Further complicating matters, organisations may have very different definitions of what constitutes compliance. For instance, one organisation may require managers to edit their schedules to reflect actual worked time, while another may consider such editing a violation. Without specific definitions of what is noncompliant, UKG chose to use unsupervised machine learning, including clustering, to detect patterns of workforce management behavior that differ from an organisation's norms.

Mitigating unconscious bias

At UKG, we believe that mitigating the risk of bias in workforce management decision making (such as employee scheduling, which might be based on availability, skill level, performance, etc.) starts with transparency. Our approach to leveraging artificial intelligence to aid in this decision making is to be very straightforward and transparent with our customers; our algorithms are tested and refined based on years of real data, but certain AI predictions are made based on the types of data fed into the analysis. To avoid so-called black box systems — where nothing is known about the inputs to the system or how it operates — we carefully evaluate the areas where AI makes most sense and are very pragmatic with respect to areas where bias could exist.

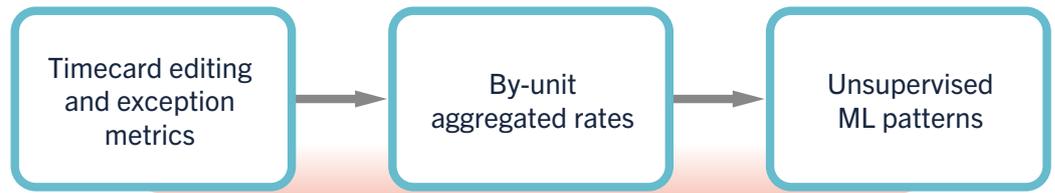


Figure 1
High-level architecture of unsupervised learning used to find timecard compliance patterns. Domain metrics such as punch-time edits are transformed into normalised (comparable across the organisation) rates at a by-unit level and then analysed for systemic patterns using an unsupervised ML technique. The end result is surfaced in dashboards and reports in the Workforce Auditor™ solution in Workforce Analytics.



Dashboards and reports

UKG innovations and advantages for compliance auditing
Figure 1 illustrates the high-level process of the UKG AI solution for detecting patterns that may be linked to compliance violations. Several key innovations are applied in this process:

- **Features** — Uncovering noncompliance patterns requires features that truly capture behavioural characteristics. UKG utilises a set of features including punch-time edits, retroactive schedule edits, and various metrics on types of exceptions in timecards. This “raw material” has been vetted not only by UKG in its AI studies but also by UKG auditors who utilised these factors in earlier manual investigations.
- **Rates for comparing units** — Locations, departments, and any other units will vary across an organisation, so one cannot simply compare the raw number of punch edits between a large department and a small department. The AI timekeeping compliance solution utilises normalised rates that allow comparisons across units and also transparency into how the comparisons were made.

- **Aggregation of patterns** — A single edit or even one week of “bad” edits may not indicate a compliance issue. The goal is instead to identify subtle, longer-term trends. Therefore, UKG uses a monthly aggregation of metric rates as well as a long-term history (usually a year) to analyse for trends.

The overall advantages of this system include:

- Organisationwide analysis using a true big-data approach to compliance analysis.
- Battle-tested metrics and normalised rates that indicate noncompliance and can be compared across units.

- Patterns detected by the AI algorithm that provide both an overview of abnormal behaviors happening in an organisation and targeted information about what units are causing them. No longer is it necessary to sift through millions of edits trying to find a needle in a haystack.

Testing and delivery

The timecard compliance solution described above has been tested on UKG customer partners across verticals and eventually was packaged into the Workforce Auditor component of Workforce Analytics™. The testing and partnerships with different customers revealed important patterns and use cases, including:

- Managers subtly manipulating punches (changing only a few minutes here and there), resulting in reduced penalties or overtime
- Managers manipulating schedules after the fact (to cover up for employees who leave early or late) to avoid organisational penalties
- Pay code edits performed by certain managers who were entering incorrect pay annotations on timecards due to lack of training

To deliver a true solution with AI pattern detection, Workforce Auditor provides both a high-level view of the patterns occurring at an organisation and drill-down capability. The latter allows users to see units within a particular pattern, how systemic the pattern is, and which subunits (or even employees) are contributing to the high numbers. Providing this actionable and targeted information to decision makers gives them a “magnet” for picking out those needles in the auditing haystack.

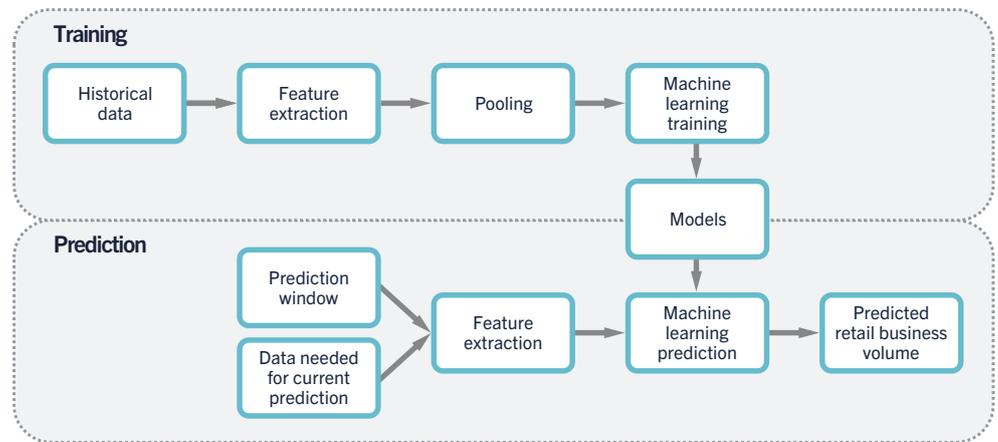
Business volume forecasting

Volume forecasting — predicting how much volume the business will process that day — is a critical first step in the process of workforce scheduling. Static formulas have been used to answer this question for over a decade, but supervised machine learning regression methods, including deep neural nets, tree ensembles, and kernel-based methods, hold the promise of higher accuracy with less manual tuning and better extensibility. UKG is delivering on this promise with the Workforce Dimensions ML forecaster, which combines a supervised ML technique with essential UKG innovations.

UKG innovations and advantages for forecasting

The architecture of the ML forecaster is shown in Figure 2. The top row of the diagram contains the training process, where the models are fitted to an individual customer’s business volume patterns based on its historical data. The bottom row shows the models used to make predictions about volumes on future days.

Figure 2
ML volume forecasting architecture. The top row covers the training process where the models are built from historical data, while the bottom row covers the prediction of individual days.



Several key UKG **innovations** in the overall process are crucial for the success of the ML forecaster:

- Features** — The use of specific features, such as recent averages, long-term trends, calendar features, organisational structure data, special events, and location data, allows the model to uncover complicated seasonal, weekly, or organisational patterns. No matter what ML algorithm is used, if the features fed to it are not informative, the predictions will not be accurate. To put it succinctly: garbage in, garbage out. The UKG engineered features used in the Workforce Dimensions ML forecaster have been proven to drive business volume and capture critical trends for real and diverse retailers.
- Pooling** — A “pooling” technique combines data from different locations across an organisation to leverage similar stores and departments when making predictions. This truly big-data approach, made possible with the UKG D5 architecture, is a complete paradigm shift from the “every department is an island” view in static formulas and brings the full power of an organisation’s data to bear on each prediction.
- Range** — The UKG machine learning forecaster is trained to accurately predict both high-volume drivers (e.g., total store sales) and low-volume drivers (e.g., refrigerators sold).
- Horizon** — Several innovations enable predicting many weeks out — a horizon that typical academic approaches do not attempt but that customers need in order to produce timely schedules.

Combining these innovations with a powerful supervised ML regression technique, the ML forecasting approach has many advantages:

- The new machine learning method shows an average improvement of over 20 percent against the static formulas across a wide variety of retailers (see below for details). The features, pooling, and range innovations all make large contributions to the accuracy gains.
- Compared with traditional static formulas, the learned model is less sensitive to outliers. This is partly due to the features used and to big-data pooling.
- Because the model is automatically fitted to the data and because of the generality of the features, very little (if any) manual configuration is needed.
- Unlike static formulas, the machine learning method is extendable; incorporating new features that impact volume does not require manually rewriting a static formula.

- By retraining the model at automatic intervals, the machine learning model can adapt over time to better fit customer data or emerging trends without manual intervention.

Testing and delivery

The advantages listed above are the result of careful testing on real-world volume data for many retailers that partnered with UKG to study the benefits of ML forecasting. Examples of results of the study for five of these customers are shown in Table 1. The two columns on the right describe the percentage of departments with forecasting improvement at each organisation and the median percentage improvement at each of those locations. These results show that for well over 90 percent of departments at all the retailers, the ML forecaster demonstrated improvements over current forecasters. The median magnitudes of the improvements are almost all over 20 percent, indicating substantial accuracy gains that can be translated into savings due to more accurate labour allocation, lower overtime rates, and other results of more precise workforce scheduling. Similar results were seen at most of these customers when comparing the ML forecaster with hand-edited forecasts.

TABLE 1: Overall accuracy improvements for the machine learning approach compared with current static forecasters.

Customer	% Departments improved (vs. current forecaster)	% Improvement (vs. current forecaster)
Restaurant	92%	23%
Supermarket 1	90%	22%
Big Box 1 (Specialty)	98%	24%
Big Box 2	90%	14%
Supermarket 2	94%	26%

The new machine learning method shows an average improvement of over 20% against the static formulas across a wide variety of retailers. The features, pooling, and range innovations all make large contributions to the accuracy gains.

Beyond research testing for accuracy, UKG has engaged a wide range of retailers to understand how best to deliver ML forecasting results. These delivery mechanisms include new interfaces and diagnostic tools for troubleshooting data issues that cause errors in a forecasting system. In addition, UKG is actively researching methods to provide transparency into “why” the ML forecaster makes a certain prediction on a given day. This information not only is helpful in troubleshooting scenarios but also provides confidence for users who rely on accurate volume forecasts to drive down-stream processes such as optimised scheduling.

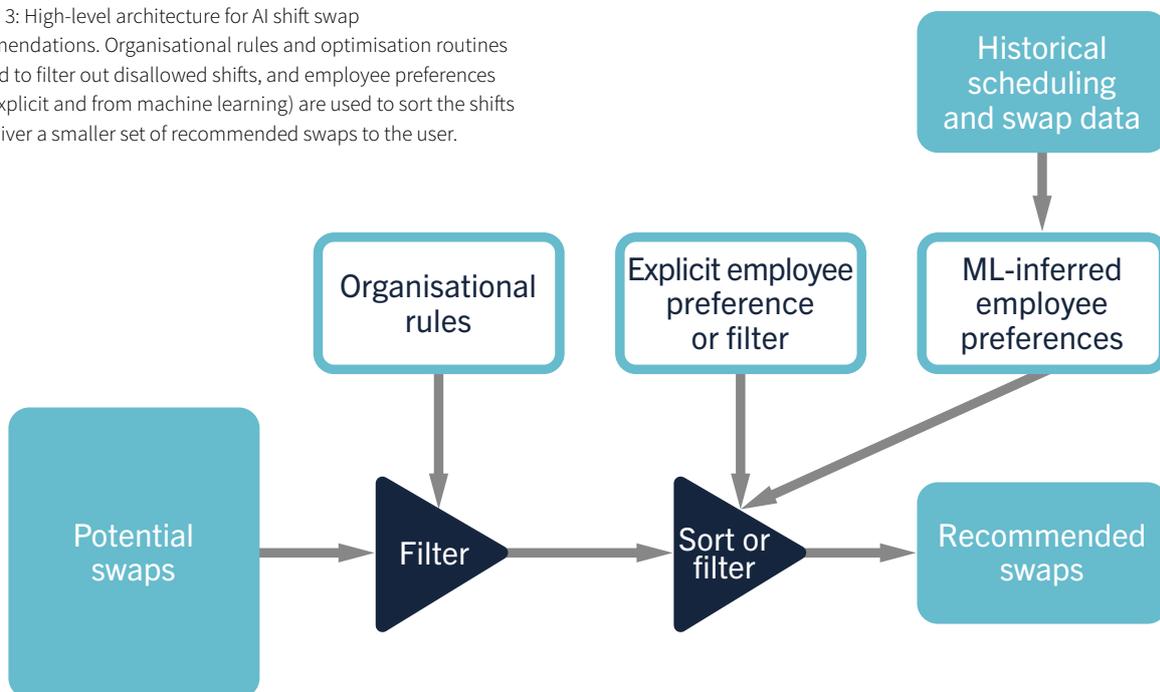
Shift swap recommendations: An emerging use case

Empowering employees to take control of their own schedules is a growing trend across industries. A common (and somewhat tedious) use case for such employees is the shift swap, when an employee is scheduled to work

a particular shift but needs to “swap” that shift with another employee’s shift on a different date and/or time. Unfortunately, that employee may have hundreds or even thousands of shifts to choose from. UKG user experience testing indicates that many swappers end up attempting multiple swaps that don’t fit rules or their preferences, leading to canceled swaps and repeated workflows.

Given these inefficiencies, a set of recommended swaps, which enforce organisational rules and account for employee preferences, can save significant employee time and frustration. Figure 3 depicts such an AI solution, which combines an optimisation routine that checks for rule violations, overtime, and other observable consequences with a machine learning approach that considers an employee’s preferences.

FIGURE 3: High-level architecture for AI shift swap recommendations. Organisational rules and optimisation routines are used to filter out disallowed shifts, and employee preferences (both explicit and from machine learning) are used to sort the shifts and deliver a smaller set of recommended swaps to the user.



UKG innovations, testing, and delivery

Once again, key innovations in the workforce management space are needed. Some of the lessons learned by UKG through user interviews and testing include:

- **Rules matter** — Organisations clearly have rules about which shifts can be swapped. For instance, at one organisation, over 93 percent of swaps were between shifts of the same length, indicating rules are in place to control overtime and other outcomes. Explicitly codifying these rules can prune swaps that violate business rules, such as excessive overtime.
- **There will always be new employees** — The UKG investigation showed that new employees are constantly entering the system. Rules and explicit preferences are crucial for making recommendations to these new employees who have no history with the organisation.
- **Employees do have (sometimes hidden) preferences** — User interviews and analysis of historical swaps indicate users have definite preferences about the people, days, and times they swap with, though these preferences are not always obvious. If Bob swaps a Tuesday shift with Sally’s Monday shift, is it because he likes swapping with Sally or prefers Monday shifts, or is there some other factor? Inferring these “hidden” reasons from examples is typically referred to in the ML community as preference learning.⁵ Mining such preferences unlocks predictive power — for instance, a prototype ML approach tested at UKG on a real customer’s data uncovered the person who would be swapped with 31 percent of the time and the day she would swap to 44 percent of the time — far better than chance.

For **delivery** of shift swap recommendations, UKG is planning a new interface and workflow for swapping in Workforce Dimensions. This interface will allow workers to easily browse or filter available shifts that do not violate scheduling rules and are ordered based on their (explicit and implicit) preferences.

AI for both management and workers

Balancing AI benefits in the workplace among managers and workers is an important and emerging topic. A key property of the business cases covered above is that all of them can show benefits to both employees and managers. For instance, the timecard compliance solution can certainly point out employees who are violating best practices. However, most of the cases UKG identified in customer data were those where a manager was manipulating timecards in management’s favour. Similarly, the shift swap recommender can save both managers and employees considerable time by stream-lining a repetitive administrative task. These cases show AI helping both workers and management by making the workplace more efficient; they illustrate how UKG uses AI technology in diverse applications with many different kinds of users.

⁵ J. Fürnkranz and E. Hüllermeier, *Preference Learning: An Introduction*, Springer-Verlag (2010).

Conclusions

This document has identified the UKG approach to AI and ML in workforce management. This approach includes UKG AI best practices: match specific AI solutions to business problems, innovate by combining AI and UKG knowledge, test models on real data, and deliver understandable and actionable results.

In addition, UKG AI advantages were covered: handling enterprisewide big data with the UKG D5 architecture, exploiting UKG knowledge, and testing with real data. The case studies highlighted these advantages and showed how unsupervised ML streamlined timecard compliance audits, how supervised ML improved the accuracy of volume forecasts by 20 percent, and how optimisation and preference learning improved the employee experience around shift swapping.

Artificial intelligence and machine learning hold great promise for streamlining work-force management processes, but there is no all-encompassing “magic algorithm.” Instead, as shown throughout this document, AI offers both toolkits that can be matched with particular business problems and knowledge that can unlock the potential strewn across an organisation’s big data. UKG and Workforce Dimensions are leading the effort to apply these technologies the right way and to continue finding new areas of application throughout the complex world of workforce management.



Building on 70 years of experience from two leaders in HR solutions, UKG™ combines the strength and innovation of Ultimate Software and UKG®. Individually, we've always put people at the center of everything we do. Together, we're committed to inspiring workforces and businesses around the world, helping to pave the way forward for our people, customers, and industry.

© 2021 UKG Inc. All rights reserved. For a full list of UKG trademarks, please visit www.ukg.com/trademarks. All other trademarks, if any, are property of their respective owners. All specifications are subject to change. SV0296-UKv2.