RAND

Research Report

AVERY CALKINS, MONIQUE GRAHAM, CLAUDE MESSAN SETODJI,
DAVID SCHULKER, MATTHEW WALSH

Machine Learning-Enabled Recommendations for the Air Force Officer Assignment System

Volume 5

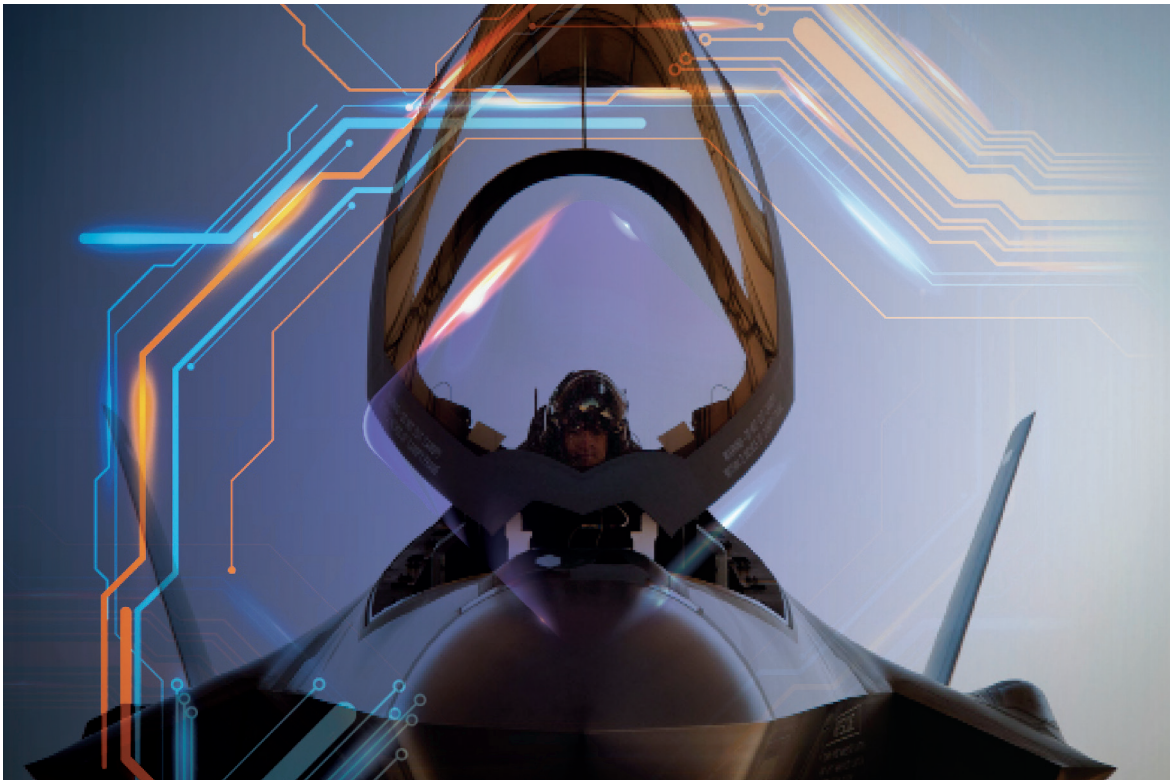

The Department of the Air Force (DAF) continues to undergo a digital transformation involving numerous efforts designed to overhaul digital architectures, systems, and processes. These wide-ranging efforts include the human resource management (HRM) domain. Specifically, the officer assignment system, which seeks "to assign the right officer to the right position at the right time to meet AF mission requirements," is undergoing a bold transformation (DAF, 2021). Along with its sister services, the U.S. Air Force (USAF) has shifted from Industrial Age, top-down personnel management procedures to a model that is closer to a digital talent marketplace (Department of the Navy, U.S. Marine Corps, 2021). This shift means that officers may apply for any available position and advertise special skills to position owners. Position owners, in turn, may review information about officers who are eligible for assignment. A shift to a true talent marketplace model, in which these preferences are the major determinant of assignments, could be a win-win, reducing the oversight burden while delivering better and more-transparent assignment matches. In addition, by reducing oversight burden, a true talent marketplace could enable more transformational change, such as allowing the USAF to move off a regular assignment cycle model.

However, implementing a marketplace approach to officer assignments comes with certain trade-offs that may limit their functionality from the point of view of the USAF. One concern is the ability of officers to discern how each assignment will contribute to their development. Another concern is whether position owners have the bandwidth to vet a potentially longer list of candidates during each assignment cycle if the system encourages more people to apply. Plus, there is no guarantee that the individual preferences of officers and position owners will lead to an outcome that is in the long-term interests of the organization, or that a talent marketplace will do a better job than the legacy process at filling less-desirable positions that are still critically important for mission effectiveness. Finally, the USAF officer assignment system is limited by the need to meet operational

KEY FINDINGS

With the advent of the Talent Marketplace, the USAF is well-positioned to implement machine learning (ML)–enhanced assignment recommendations. Besides providing a platform to deliver recommendations, the Talent Marketplace can be used to gather additional information about officer qualifications, position details, and historic preferences, which would improve the effectiveness of future recommendations.

- **A useful way to implement ML-enhanced recommendations would be to organize them around the themes of officer development, performance, and job satisfaction because they capture the factors most likely to be important to officers, position owners, and the career field assignment team.** Recommendations organized around these themes are also easier to construct and implement than a single holistic recommendation for each officer.
- **ML-enabled recommendations augment, rather than automate, human decisionmaking.** Our suggested design includes many points where humans provide key inputs to inform ML recommendations. The virtue of a recommendation system is that it allows officers, position owners, and assignment teams to flexibly apply additional knowledge beyond the scope of the ML model.
- **For most officers, position owners, and assignment teams, the number of options available in a single assignment cycle might not be large enough to necessitate personalized recommendations.** However, ML-enabled recommendations are an enabling technology to facilitate a transition to a more flexible assignment system that selects officers for rotation according to development needs rather than assigning a set of officers who have been preselected for rotation.

Abbreviations

ADP	Airman Development Plan
AI	artificial intelligence
DAF	Department of the Air Force
HRM	human resource management
KSAs	knowledge, skills, and abilities
ML	machine learning
NLP	natural language processing
OPR	officer performance report
USAF	U.S. Air Force

requirements within the constraints on assignments that officers can hold (e.g., due to disqualification or family medical needs)—meaning that the officer assignment system cannot function as a true free market. Because of these concerns, the Air Force Talent Marketplace does not yet function as a *true* talent marketplace. Instead, it is a tool for officers and position owners to communicate preferences to assignment teams, who then make officer assignments in a process like the legacy system that it was supposed to replace.

Marketplaces have been transformed by the digital age, and all the detailed information they contain is archived for posterity. This creates opportunities to improve the functionality of marketplaces by leveraging digital information in artificial intelligence (AI)– and machine learning (ML)–enabled decision-support tools for future users, which could alleviate some of the concerns about career development, bandwidth, and organizational interests. An obvious way to start improving marketplace decisions is to *consider incorporating recommendation engines*—systems that assist with exploring large databases by presenting or ranking options according to the user's needs or preferences. Incorporating ML-enabled recommendations could enable the transition from the legacy USAF Officer Assignment System to a genuine marketplace model by helping to align preferences of officers and position owners with the USAF's priorities. Even in the absence of a shift to a true talent marketplace model, recommendations could improve the experience of officers and position owners by helping them find good potential matches that they might not have otherwise considered.

What would change following the adoption of a recommendation engine? To return to the limitation examples above, recommendation engines could suggest positions to officers that will allow them to develop new competencies. Likewise, recommendation engines could provide an ordered list of candidates to position owners based on alignment between candidates' knowledge, skills, and abilities (KSAs) and position needs. These and other benefits to

RECOMMENDATIONS

- The USAF should experiment with using a matching algorithm to combine officer and position owner preferences. As implemented, the Talent Marketplace makes information visible and increases the transparency of assignments. Using a matching algorithm is the next step toward transitioning to a true, decentralized marketplace that matches officers to positions based on mutual preference—an approach that has been shown to be beneficial in other markets. However, the matching algorithm must consider the need to meet warfighting requirements and to accommodate special circumstances, such as by-name requests and family medical needs, and the assignment team members should be transparent about the process they use and the constraints they face should the USAF shift to an algorithmic approach to assignments.
- The USAF should experiment with delivering position and candidate recommendations using recommendation engines. This may enhance the current Talent Marketplace, and it may reduce risks associated with transitioning to a true, decentralized marketplace.
- The USAF should gather new sources of data, such as officers' satisfaction with positions and position owners' satisfaction with officer performance. Aside from directly benefiting the future performance of recommendation engines, this information can provide valuable information to other USAF HRM processes.

implementing recommendation engines in the Talent Marketplace system can be summarized as follows:

- **Reducing the workload of evaluating candidates and positions.** By culling the applicant or positions list, a recommendation system could save position owners and officers time.
- **Improving the quality of assignment matches.** By incorporating richer information on officer and position owner preferences, data-driven recommendations could help each party provide better inputs into the matching process.
- **Enabling more transformational change.** Recommendation systems are ideally suited for exploring large databases intelligently. By making the officer-assignment matching process easier, a recommendation system could allow the USAF to expand the number of officers and positions under consideration. This would allow the transition from an assignment process that matches only officers who have been preselected for rotation to one that continuously considers all officers and positions to best meet institutional and developmental needs.

Notwithstanding the potential benefits of using a recommendation system with the Talent Marketplace, the design and evaluation requirements for a recommendation system for the USAF will, by necessity, differ from those for other consumer markets. The singular goal of Netflix recommendations, for example, is to show users content that they will like. By contrast, the challenge of recommending assignments that both meet warfighting requirements and balance development needs and individual desires over 20- to 30-year careers requires careful forethought of additional dimensions.

Organization of the Report

The goal of this report is to assess the affordances of USAF digital architectures, systems, and processes, including the newly introduced Talent Marketplace, to deliver machine-generated assignment recommendations to officers, position owners, and assignment teams. In addition to considering the feasibility and utility of doing so, we identified potential pitfalls for future users of the USAF Talent Marketplace. The remainder of this report is organized as follows:

- We describe types of recommendation systems commonly used in labor markets.

- We detail officer assignment systems and processes as they exist today.
- We present a proposal for how machine-generated recommendations could be incorporated into assignment systems and processes.
- We close by summarizing the main findings and offering recommendations.

Recommendation Systems in Labor Markets and Their Applicability to the Talent Marketplace

In this section, we describe types of recommendation systems used in labor markets and link them to the Talent Marketplace. We begin by describing approaches used in civilian labor markets. We then describe officer assignment systems and processes as they exist today. We end by commenting on the applicability of these approaches to officer assignment.

Types of Recommendation Systems

Recommendation systems present options (e.g., products or, in this case, positions or candidates) according to an underlying model of what the user wants. There are three general approaches to generating recommendations: *knowledge-based*, *content-based*, and *collaborative filtering*. Figure 1 describes these types of systems, as applied to officers who are considering assignments (Al-Otaibi and Mourad, 2012). *Knowledge-based approaches* generate recommendations based on expert-encoded rules. *Content-based filtering* constructs a model of the user's preferences based on their inputs or on their past likes and dislikes and recommends options according to those preferences. Finally, *collaborative filtering* predicts the unknown preferences of a new user based on the known preferences of a community of similar users. Furthermore, *hybrid systems* layer elements of these and other approaches together, yielding a wide variety of design possibilities for systems to generate recommendations.

The chief benefits of knowledge-based recommendations are twofold. First, they do not require a large amount of existing data—the rules used to generate recommendations are explicitly coded based on expert knowledge. Second, recommendations are highly interpretable. The main drawback is that the knowledge engineering required to build and maintain large systems may be cost prohibitive (Kendal and Creen, 2007). In addition, knowledge-based recommendations may maintain the status quo without yielding new insights.

The benefits and drawbacks of content-based and collaborative-filtering approaches are the opposite of those of knowledge-based approaches. They do

FIGURE 1
Types of Recommendation Engines

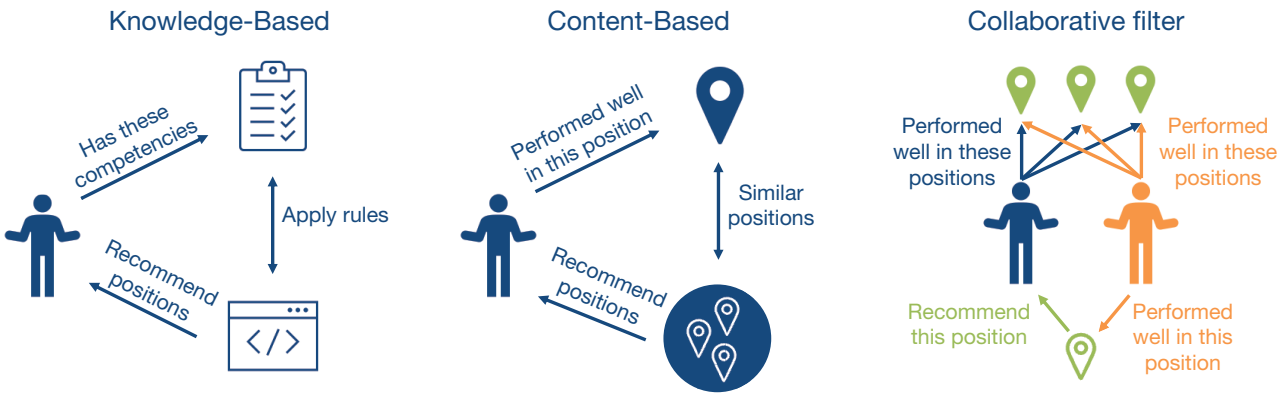

not require domain knowledge to construct explicit rules—the preference models used to generate recommendations are learned. The quality of the models improves over time as the amount of historical data increases. In addition, learned models may lead to fortuitous recommendations that the user would not otherwise have considered. The main drawback of these approaches is that they require a large amount of data to learn valid models of user preferences.

Variations of these methods are used in the private-sector labor market. For example, knowledge-based systems map key words in the job seeker's resume to a dictionary of terms (i.e., an ontology) and present positions that match the individual's KSAs. Content-based systems present positions like those the job seeker has applied for in the past. Finally, collaborative-filtering systems present positions that a community of users like the job seeker have applied for in the past. Conversely, these systems are also used to recommend candidates to organizations seeking to fill positions.

USAF assignments have much in common with the civilian labor market, but the USAF assignment system differs in three fundamental ways. First, assignments must meet immediate USAF needs, but they serve the additional purpose of developing officers for future leadership roles and responsibilities. Civilian firms must also develop employees, but they have the option to hire external candidates with the required skills and abilities at every level.

Second, the supply of jobs and candidates in the civilian labor market is effectively endless. In the USAF assignment system, positions and candidates are finite and zero-sum; during each assignment cycle, every available officer must be matched with a position, and an opportunity gained by one officer is an opportunity lost by another. The closed nature of the USAF assignment system resembles a matching market. In fact, the National Residency Match, a solution for matching medical students to residency programs, was used as a starting point for generating assignment matches in the original implementation of the Talent Marketplace (National Academies of Sciences, Engineering, and Medicine, 2021).

Third, in civilian labor markets, individuals are free to pick their employers. Officers have less flexibility in picking positions. After receiving an assignment, they must accept the assignment or separate from the USAF (i.e., *seven-day opting*). Thus, the USAF has greater ability to fill less-desirable positions.

The supply of jobs and candidates in the civilian labor market is effectively endless.

Next, we review the structure of the Air Force Talent Marketplace and the preferences and needs of stakeholders involved in the assignment process. We then describe how recommendation systems can be incorporated into the Talent Marketplace and the data requirements to do so.

Structure of the Officer Assignment System

The USAF Talent Marketplace is a web-based platform that allows officers in the grades of O-1 through O-5 to express assignment preferences. The intent of the Talent Marketplace is to coordinate human capital management while maintaining an agile and multi-capable workforce. The Talent Marketplace is housed in an online system called MyVector, which contains notifications of vacancies and officers' profiles and is the system in which officers and position owners submit information on their preferences, among many other features.

Figure 2 gives a high-level view of how the USAF Talent Marketplace functions, including the potential role of recommendations. First, officers eligible for an assignment are notified. This gives officers an opportunity to self-advertise their capabilities while enabling leadership to provide input to their career progression. At the same time, the job owners (i.e., billet owners or position owners) advertise vacant positions. Second, officers review information about vacant positions and submit ordered lists of assignment preferences. Third, job owners review informa-

FIGURE 2

High-Level View of the Talent Marketplace

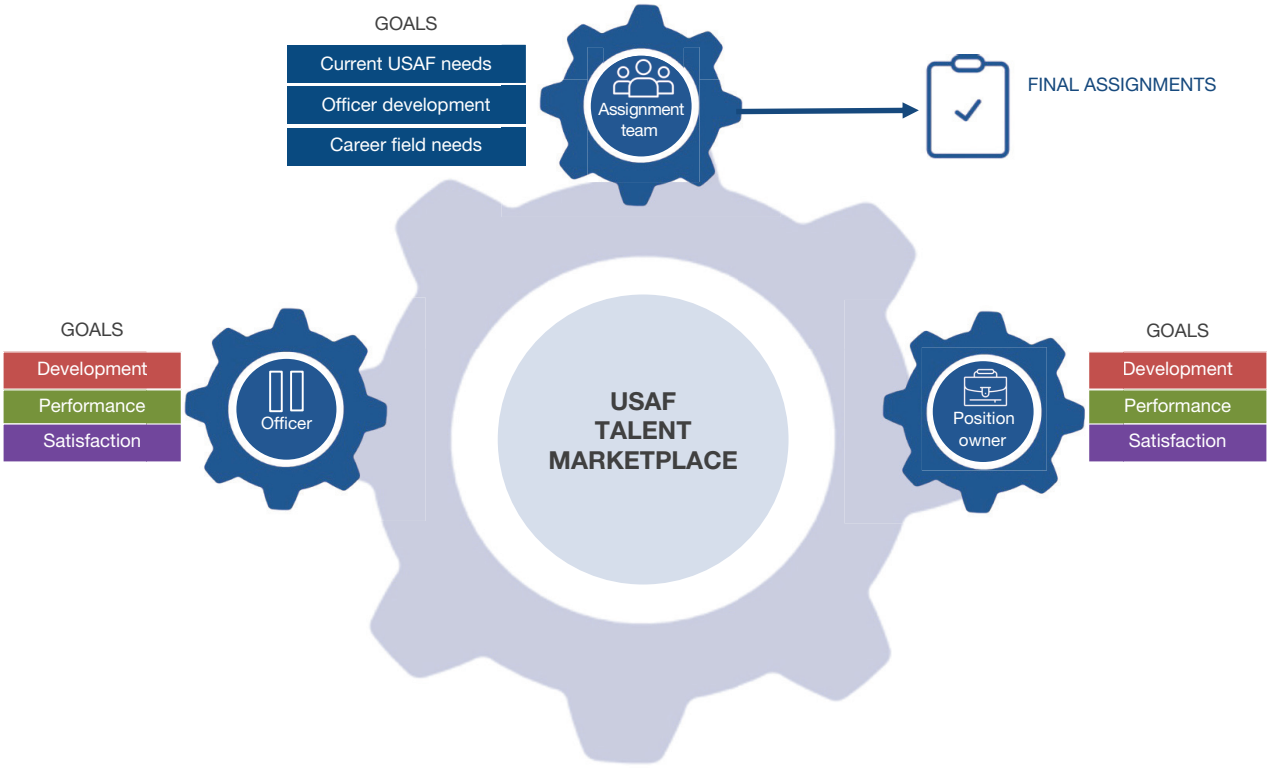

tion about interested candidates (i.e., billet volunteers) and submit ordered lists of preferred candidates (i.e., bids, which can also include officers who did not volunteer for the position). Fourth, career field managers and leadership provide additional input to help assignment teams balance the professional needs and preferences of individuals with the institutional needs of the USAF. The assignment team reviews the preference information and matches officers to positions. Finally, after matches are reviewed, officers are notified of their new assignments.

Stakeholders in the Officer Assignment System

The three primary stakeholders in the officer assignment system are officers, position owners, and assignment teams. Members of each of these groups have their own goals, which has implications for the structures of the recommendation engines used for them.

Officers

The officer's goal is, quite simply, to get the best possible position. Yet this is surprisingly complex. The best possible position depends on an officer's career field and military and life goals, and different officers may use different processes to decide whether a job is a good match.

In standard economic theory, an individual's preferences for a differentiated product (e.g., one car model versus another) are modeled based on the characteristics of that good and how much the individual cares about each of those characteristics. That is, a product is treated as, essentially, a collection of its characteristics rather than as a whole item, and consumers' preferences over a set of differentiated products are based on how much they value the characteristics of the item (McFadden, 1973). This theory can be adapted to a labor market in which individuals are choosing different jobs: Their preferences for that job will depend on the job's characteristics and

how much the individual weights those characteristics in their preferences.

In the case of assignments, officers may consider such position characteristics as mission, location, work environment, and development value. The combination of characteristics partially determines how interested the officer is in one position relative to others. Additionally, each officer may place different weight on each of the characteristics. For example, some officers may care a great deal about work environment, while others may care about career advancement. Recommendations must reflect the characteristics of different positions and how much the individual cares about each characteristic.

Position Owners

The position owner's goal is to get the most qualified officer. This depends on the characteristics of the officer and the position. From a modeling standpoint, performance is itself a reflection of the quality of the match between the officer and position. The recommendation engine should consider these characteristics when presenting candidates to the position owner; one way to do that is to consider the characteristics and performance of individuals who previously worked in the position.

Assignment Team

Assignment teams are ultimately responsible for assignment decisions. They must consider inputs from career field managers and major command functional managers. As noted in the introduction, the assignment team's first goal is "to assign the right officer to the right position at the right time to meet AF mission requirements" (DAF, 2021). The assignment team must also consider the professional development of the officers, along with commander recommendations, officer preferences, position owner preferences, and other constraints, such as disqualification codes, by-name requests, family medical needs, and manning priorities (DAF, 2021).

Officer and position owner preferences are currently sent to assignment teams via the Talent Marketplace. Meeting these preferences is important to reduce dissatisfaction, thereby increasing retention, performance, and readiness.

The stakeholders in the officer assignment system have three major goals beyond meeting valid warfighting requirements: officers' development, performance, and job satisfaction.

Altogether, the stakeholders in the officer assignment system have three major goals beyond meeting valid warfighting requirements: officers' development, performance, and job satisfaction. Though each party values each goal differently, each party has reason to consider each factor when determining the best set of officer-position matches.

An Approach to Implementing Recommendations in the Talent Marketplace

The USAF assignment system resembles labor markets in some ways. The officer is like a job seeker, and the position owner is like the employer. Thus, knowledge-based systems can recommend positions that an officer is qualified for, content-based systems can recommend positions based on a model of the officer, and collaborative filters can recommend positions based on other similar officers. Likewise, knowledge-based systems can recommend officers who meet the qualifications for a position, content-based systems can recommend officers based on a model of the position, and collaborative filters can

recommend officers based on other positions like the one being filled.

As these examples illustrate, it is conceptually possible to apply recommendation systems to officer assignments. Further, as described in Volume 2 of this series (Walsh et al., 2024), which presented a framework to prioritize ML projects in the HRM domain, applying ML to officer assignments directly addresses the HRM objectives of improving individuals' performance and of enhancing opportunities available to them—and, furthermore, it can indirectly address the objectives of improving business processes and of increasing individuals' motivation.

Yet officer assignments differ from civilian labor markets in key ways, as described in the introduction. Mainly, the assignment system is a closed system—all officers must be assigned to positions, and most positions must be assigned an officer. This may create conflicts in terms of what is best for individual officers and what is best for the workforce. Additionally, assignments must be forward-looking to provide officers with experiences that will prepare them for still greater responsibility in future positions. These considerations complicate but do not preclude the use of recommendation systems for officer assignments, as we discuss in the next section.

Overview of Proposed Recommendation System

This section outlines potential approaches to constructing recommendation systems for officer assignments. We begin by describing a functional design of recommendations for officers and position owners. Our proposed design gives three separate streams of recommendations for officer-position pairing based on (1) developmental value, (2) performance, and (3) satisfaction, and it provides a starting-point set of officer-position pairs to the assignment team. We then describe how officers and position owners can use these recommendations along with other information to reach decisions. Finally, we describe how assignment teams could use officers' and position owners' preferences and recommendations as an aid to creating a set of officer assignments, given that the assignment team faces a different set of needs and constraints from those faced by officers and position owners.

Thematic Recommendations

The motivations of officers and position owners reflect different combinations of three common themes:

1. *Developmental value*: the extent to which a position will provide an officer with KSAs to prepare them for future leadership roles and responsibilities.
2. *Performance*: the extent to which an officer's existing KSAs will allow them to meet position demands.
3. *Satisfaction*: The extent to which an officer will be satisfied with a position, for reasons beyond developmental value and performance.[1]

We propose generating general streams of recommendations around these themes and providing them to officers and position owners, rather than building holistic streams of recommendations for the different stakeholders, as shown in Figure 3. This type of thematic organization is attractive for several reasons. First, both officers and position owners care about these themes, but the relative importance of each varies by group. Representing the themes separately eliminates the need to learn how each group and individual weights the themes. Second, thematic recommendations are easier to interpret than a holistic recommendation that combines the three. Third and finally, the USAF may currently lack data sources to adequately capture all three themes. Represented separately, each stream of recommendations can be incrementally refined using all the most relevant data available.

We also propose providing a recommendation of a set of officer-position pairs based on submitted preferences to the assignment team, who could use it as a starting point in the assignment process. The assignment team could also be provided with the recommendations shown to each officer and position owner, which might provide context for why officers or position owners submitted particular lists of preferences.

FIGURE 3

High-Level View of the Proposed Talent Marketplace

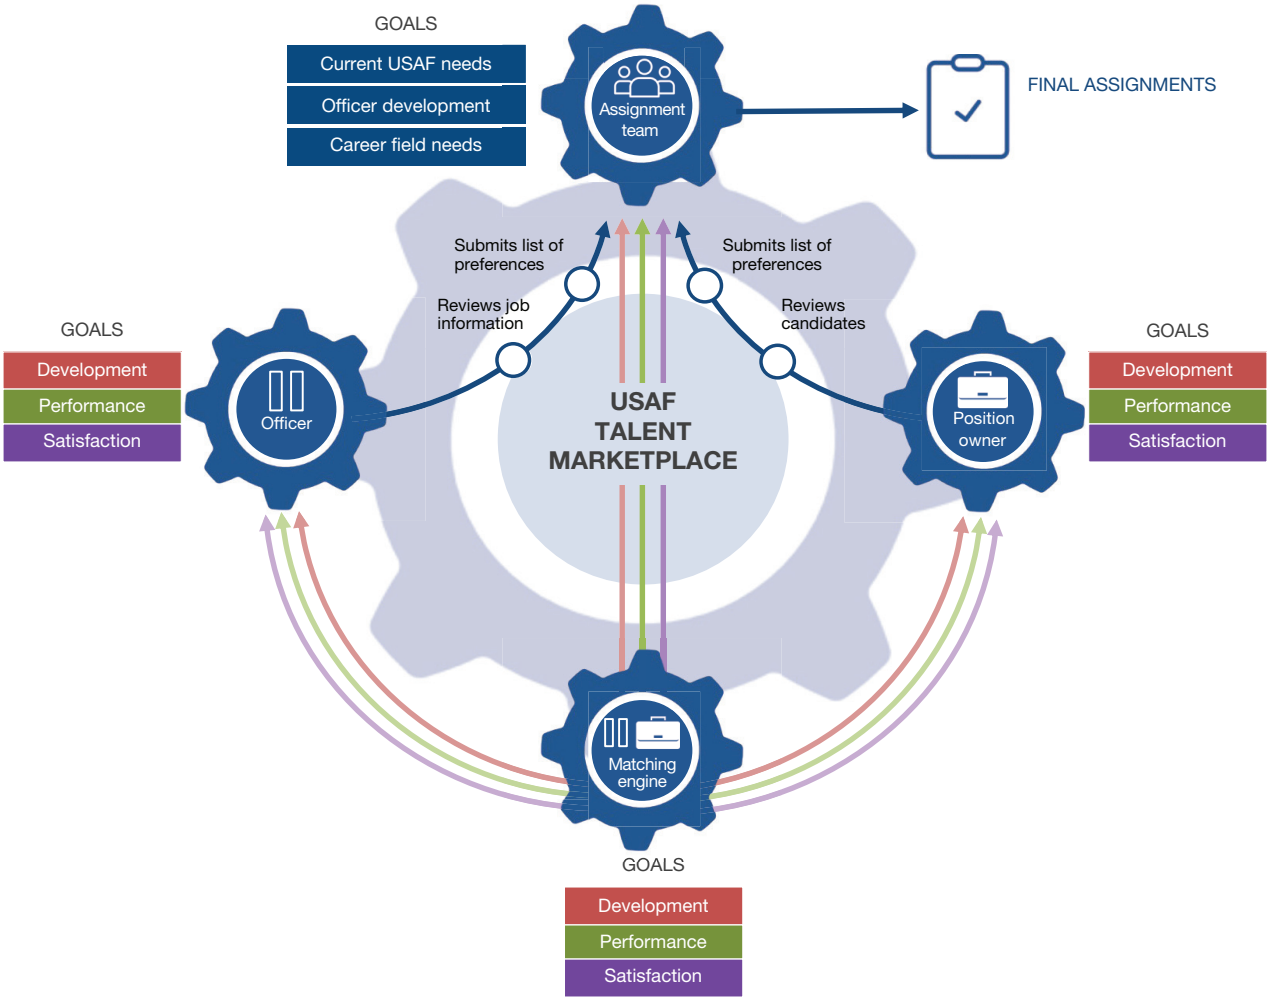

Recommendations for Officers and Position Owners

Construction of Thematic Engines

Table 1 describes how the different types of recommendation engines can be applied to the themes from the perspective of the officer seeking a position along with required inputs and, in the case of content-based and collaborative filters, the required outcomes to enable them. With slight modifications, Table 2 captures the position owner's perspective.

Knowledge-Based Approaches

Knowledge-based approaches recommend positions based on expert-defined rules. The knowledge-based approach for development recommendations uses domain knowledge of assignment teams to create a set of rules to recommend positions with high development value. Assignment teams maintain a theoretical model of development (an Officer Career Development Brief, or previously a *career pyramid*), and they use heuristics to identify the types of positions that are appropriate for officers at different points in their careers. These business rules should not be overlooked, as they have evolved to meet senior-level development needs and have stood the test of time. The knowledge-based approach codifies these rules.

The knowledge-based approach for job performance specifies KSAs needed as prerequisites for high performance in a position. These prerequisites are compared against an officer's KSAs—inferred

TABLE 1

Descriptions and Data Requirements for Officer Recommendations

	Recommendation Type			
Theme	**Knowledge-Based**	**Content-Based**	**Collaborative Filter**	**Required Data**
Development	Encode and apply rules used by assignment teams	Find positions related to development goals expressed by airmen or supervisors	Find positions that contributed to the development of similar officers	Board results such as outcomes, order of merit, or scores
Performance	Find positions with prerequisites that officer meets	Find positions like ones the officer has performed well in	Find positions that similar officers have performed well in	OPRs or position owner ratings
Satisfaction	Find positions with characteristics that the officer specifies	Find positions like ones the officer has had high satisfaction with	Find positions that similar officers have had high satisfaction with	Climate surveys or officer ratings
Required inputs	SME inputs	Personnel data	Position data	

NOTE: OPR = officer performance report; SME = subject-matter expert.

TABLE 2

Descriptions and Data Requirements for Position Owner Recommendations

	Recommendation Type			
Theme	**Knowledge-Based**	**Content-Based**	**Collaborative Filter**	**Required Data**
Development	Encode and apply rules used by assignment teams	Find officers who would benefit from development opportunities offered by position	Find officers who would benefit from development opportunities offered by similar positions	Promotion results
Performance	Find officers who meet position prerequisites	Find officers like ones who have performed well in position	Find officers who performed well in similar positions	OPRs or position owner ratings
Satisfaction	Find officers who specified characteristics of position	Find officers like ones who had high satisfaction with position	Find officers who had high satisfaction with similar positions	Climate surveys or officer ratings
Required inputs	SME inputs	Personnel data	Position data	

from assignment histories and information contained in OPRs and other sources—to exclude certain positions from consideration. Likewise, the knowledge-based approach for job satisfaction specifies position attributes, such as location. Once again, these attributes are compared against an officer's requests to exclude certain positions from consideration.

As discussed earlier, the primary benefit of the knowledge-based approach is that it does not require a large amount of data, only expert knowledge. In addition, recommendations are explainable. The primary drawback is that the quality of recommendations is limited by the quality of expert knowledge. A knowledge-based approach cannot, for example, discover new and more-effective development pathways, or positions outside of an individual's experience in which they would perform well or be highly satisfied. In addition, it may be prohibitively costly to create and maintain a knowledge-based system. For development recommendations, it may be infeasible to scale assignment rule sets across all functional areas and to maintain them as development paths and requirements change over time. For job performance and satisfaction, officers and position owners would need to supply complete and accurate information to the Talent Marketplace.

Content-Based Approaches

Content-based approaches recommend positions based on a user profile. These approaches require detailed position data to compute similarity between assignments. The USAF does have administrative data on hundreds of thousands of positions spanning multiple decades. Administrative data could be augmented using free text contained in the duty description portion of OPRs and, in the future, data extracted from position descriptions in the Talent Marketplace.

The content-based approach *for development* requires an officer profile, such as an Airman Development Plan (ADP), that expresses an officer's current development needs. The recommendation engine could then identify positions that are semantically related to the types of experiences expressed in an ADP.[2]

The content-based approach *for performance* recommends positions like those in which the officer has performed well in the past. This requires quantification of job performance, which could be done in two ways. First, OPRs contain a narrative description of officers' past performance in positions. Text-mining methods such as sentiment analysis can be applied to these data to construct a numeric estimate of an officer's performance. Recent work has demonstrated the feasibility of using natural language processing (NLP) and ML to do so (Schulker, Lim, et al., 2021). For example, the Personnel Records Scoring System (PReSS) described in Volume 3 of this series was created to help board members to use information contained in OPRs to make developmental education and promotion decisions (Schulker, Williams, et al., 2024). PReSS could be extended to support assignment decisions. Second, position owners could rate the performance of outgoing officers in a future version of the Talent Marketplace.

The content-based approach *for satisfaction* recommends positions like those in which the officer had high satisfaction in the past. This requires quantification of satisfaction, which could be done in two ways. First, career outcomes, such as separation, could indirectly reflect satisfaction. Second, outgoing officers could rate their satisfaction with positions in a future version of the Talent Marketplace. The USAF does not systematically collect information on job satisfaction and associated position characteristics (e.g., work-life balance, stress, local nonwork amenities; Schulker, Harrington, et al., 2021). These data would be needed to enable satisfaction-based recommendations.[3]

Collaborative-Filter Approaches

Collaborative filters base recommendations on a community of individuals who are like the officer; therefore, these approaches require detailed personnel data to compute similarity between officers. As with the content-based approach, a collaborative-filter approach could draw on the USAF's administrative data (covering hundreds of thousands of individuals over multiple decades) augmented by free text from OPRs and, in the future, data extracted from officer profiles in the Talent Marketplace.

The collaborative-filter approach for development recommendations identifies positions that contributed to the growth of other individuals who are like the officer. This requires quantifying development value, which could be done in two ways. First, the USAF could treat historical assignments, generated by assignment teams, as reflecting informed decisions. In this case, the collaborative filter would learn to emulate all historical assignments. Second, the USAF could focus on assignment sequences for individuals who were promoted to senior grades (e.g., O-6). In this case, the collaborative filter would learn to emulate historical assignments of successful officers.

The collaborative-filter approach for performance would recommend positions in which other individuals like the officer have performed well. Likewise, the collaborative-filter approach for satisfaction would recommend positions in which other individuals like the officer have had high satisfaction.

The primary benefit of the content-based and collaborative-filter approaches is that they do not require domain knowledge to construct an explicit set of rules, only data about historical assignments and their outcomes. In addition, by learning from historical data, these approaches may come to produce more-effective development pathways, improve job performance, and increase satisfaction. The primary limiting factor is that these methods require a

large amount of historical data, along with ways to assess the quality of historic outcomes.

Extension of Approaches to Position Owners

Performance and satisfaction recommendation streams can also be generated for position owners. For example, content-based approaches can be used to identify officers like those who previously performed well in or had high satisfaction with the position. On the other hand, collaborative-filter approaches can identify officers who previously performed well in or had high satisfaction with similar positions. Finally, knowledge-based approaches can identify the subset of officers who meet position prerequisites or who have not screened out the position based on its attributes.

Incorporation of Thematic Engines into the Talent Marketplace

The thematic recommendations can be incorporated as a new capability into MyVector's search pages for positions to apply to (on the officer side) and officers to interview (on the position owner side). Figure 4 shows a mock-up of a potential user interface design for the officer side. The interface would remain largely the same as it currently appears, with recommendations added at the top of the page under such headings as *Greatest development value*; *Greatest fit for position*; and *Greatest satisfaction*. The headings refer to the development, performance, and satisfaction themes, respectively, and each field will display positions. Additionally, the Talent Marketplace would include a fourth stream of recommendations, titled *Best fit for you*, that combines the three streams of

FIGURE 4

Potential User Interface Design for Incorporation of Recommendations into the Talent Marketplace

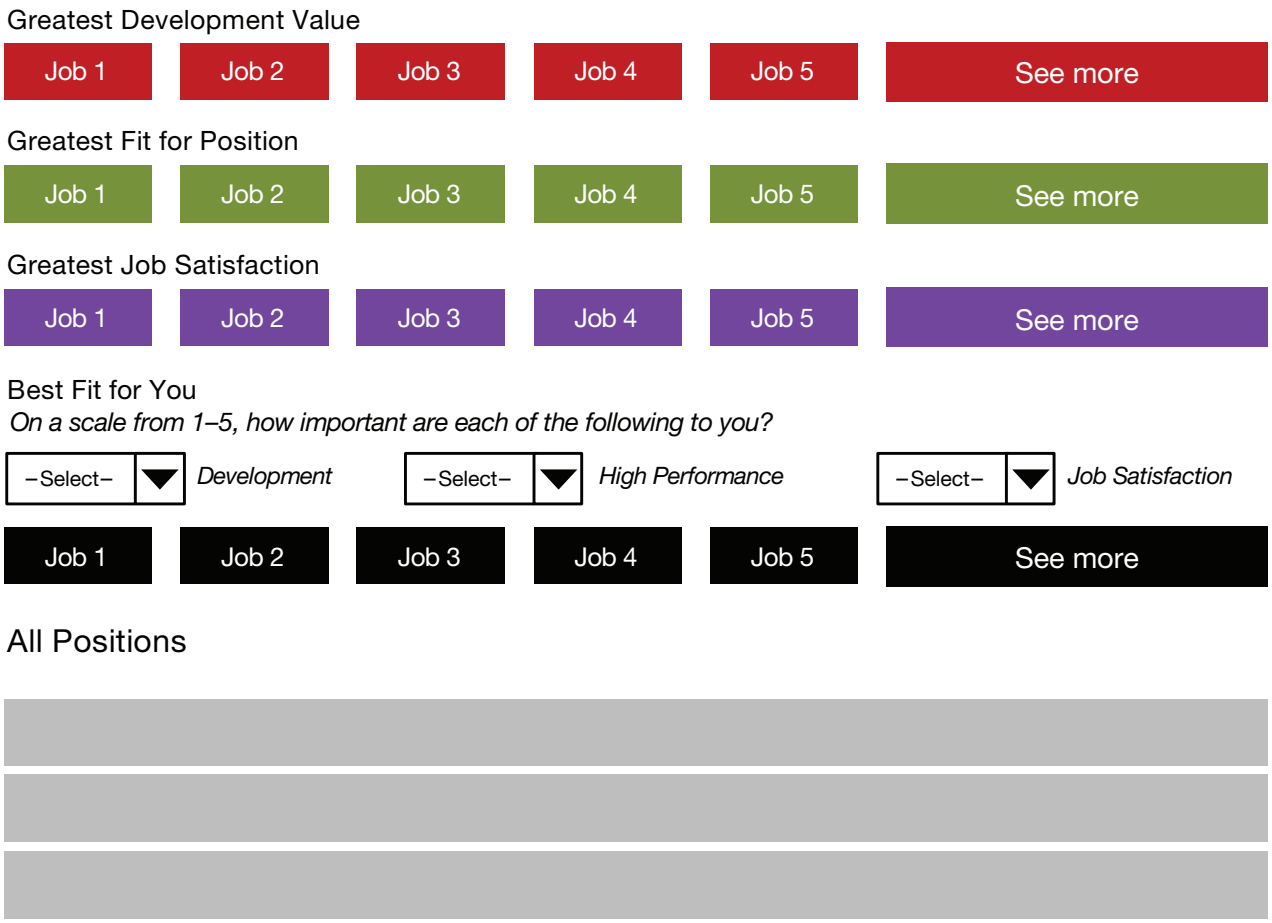

recommendations based on the relative importance of each theme to the individual. In this example, the top position recommendation is implemented by asking individuals to rate the importance of the three themes right on the page where recommendations and job postings are presented, allowing easy customization, but information could also be collected via user profiles. Officers would maintain the ability to look through all positions available to them.

An analogous user interface to that suggested in Figure 4 could be used for position owners. For position owners, a complementary set of recommendations would appear with the same headings (i.e., *Greatest development value; Greatest fit for position;* and *Greatest satisfaction*), with the fields displaying candidates. As with officers, the position owner could view a fourth stream of recommendations that combined the three thematic streams based on the relative importance of each to the position owner, either drawn from their user profile or entered on the page where officer profiles and recommendations are presented.

Aside from displaying the recommendations on the search pages, the Talent Marketplace would function largely the same way it does now. Once officers and position owners have submitted preferences, these preferences and the streams of recommendations provided to each officer and position owner would be sent to the assignment team performing the remainder of the process. The contribution of the recommendations would be the degree to which they help each stakeholder to provide more-accurate inputs into the process, thereby improving the final set of matches.

Recommendations for the Assignment Team

In each assignment cycle, officers and position owners solve a local problem—finding the position or individual that best meets their needs. The assignment team must solve a global problem—the optimal pairing of all officers eligible for assignment to positions during the assignment cycle. As outlined earlier, the assignment team also has additional goals and faces additional constraints that the officers and

In each assignment cycle, officers and position owners solve a local problem—finding the position or individual that best meets their needs.

position owners may not: The assignment team, first and foremost, must meet the needs of the Air Force and the career field and must also accommodate special circumstances, such as by-name requests from senior leaders and medical needs.

The assignment team should have access to all information, including officer and position owner preferences and the recommendations provided to officers and position owners, for its decisionmaking. The recommendation to the assignment team should provide a full set of officer-position pairs, which can be updated to incorporate changes made by the assignment team and which ensures that all critical positions are filled and that no officer is assigned to a position they are unable to take.

Several algorithms can accomplish these goals. We recommend a preference-based matching algorithm, such as the one used in the National Residency Match that pairs physicians and residencies, with modifications made to ensure that critical positions are filled and that airmen with exceptional needs are matched to suitable positions (Roth and Sotomayor, 1992; Roth and Peranson, 1999). These algorithms would match officers to positions using ranked lists of preferred positions from officers and ranked lists of preferred officers from position owners, just like those already submitted to the assignment team through the Talent Marketplace. Such an algorithm has previously been proposed by the National Academy of Sciences, Engineering, and Medicine to improve the legacy USAF Officer Assignment System

(National Academies of Sciences, Engineering, and Medicine, 2021).

A preference-based matching approach is attractive for two reasons. First, it is transparent, avoiding concerns that officers might have about the automation of assignment processes produced by a black-box ML algorithm. Second, the simplest version of the proposed algorithm has been shown to be optimal for officers—that is, officers will be assigned to their top-ranked position, given that the position owner also expressed a preference for the officer (Gale and Shapley, 1962). While that feature might no longer hold after the modifications made to ensure that critical positions are filled and that officers with exceptional needs are matched to suitable positions, the number of matches between officer and position owner when both choose the other can be tracked and maximized using a matching algorithm approach. By treating officer preferences in this way, the matching approach may improve officer retention and readiness.

The preference-based matching approach provides initial recommendations, but the assignment team is ultimately responsible for assignment decisions. Indeed, the assignment team may have additional information that officers and position owners do not, such as an equity requirement that limits who may be assigned to a short tour. It is unlikely that any matching algorithm could truly take all the assignment team's constraints into account. Thus, it is unlikely that the assignment team could ever take the matching algorithm's output as the final set of assignments without making any modifications. We therefore suggest an iterative process, whereby the assignment team can make modifications to the set of recommended officer/position pairs, such as accepting some matches that are particularly good, forbidding some matches that are particularly bad, or reassigning an officer to a different position or a position to a different officer. The assignment team could then re-run the algorithm conditional on those modifications to create an updated list of recommendations. Ideally, the iterative approach would reduce the burden on the assignment team that comes out of making changes to the set of assignments suggested by the matching algorithm, especially if changing one suggested assignment creates a domino effect where many other matches must be updated.

A final execution challenge is the possibility that officers or position owners will object if they do not receive one of their desired assignments, especially in the case where the assignment received followed a modification made by the assignment team. The USAF could mitigate this execution challenge by increasing transparency of the assignment process. That is, the matching approach and the types of modifications made by the assignment team should be communicated clearly to officers and position owners. In particular, the assignment team should openly acknowledge the concerns that would lead them to modify the assignments suggested by the matching algorithm, such as warfighting requirements, family medical needs, equity considerations, by-name requests from senior leaders, and any other special circumstances that would require a particular assignment to be made. The assignment team should also communicate how they choose officers for less-desirable positions—for instance, in the case of equity requirements for who can be assigned a short tour, the assignment team should describe the criteria for determining whose "turn" it is to be assigned a short tour.

Ideally, the recommendations to the assignment team would be presented in a web-based interactive tool. Figure 5 shows a potential user interface for such a tool. The top panel displays a summary view, listing all officer/position matches and (if applicable) all unmatched officers and unmatched positions. The bottom panel displays the view the assignment team would have if they click "see more," which would allow them to view more-detailed information about the officer and position. In this example, the tool shows the recommendation streams provided to the matched officer–position owner pair, along with the preferences expressed by both. Both views allow the assignment team to take any of the actions described in the previous paragraph, make modifications, and click the "re-run with my changes" button at the top of the summary view to see a new set of assignment recommendations.

The tool could also allow the assignment team to use one of the three streams of recommendations provided to officers and position owners in place of

FIGURE 5

Potential User Interface for Assignment Team Tool

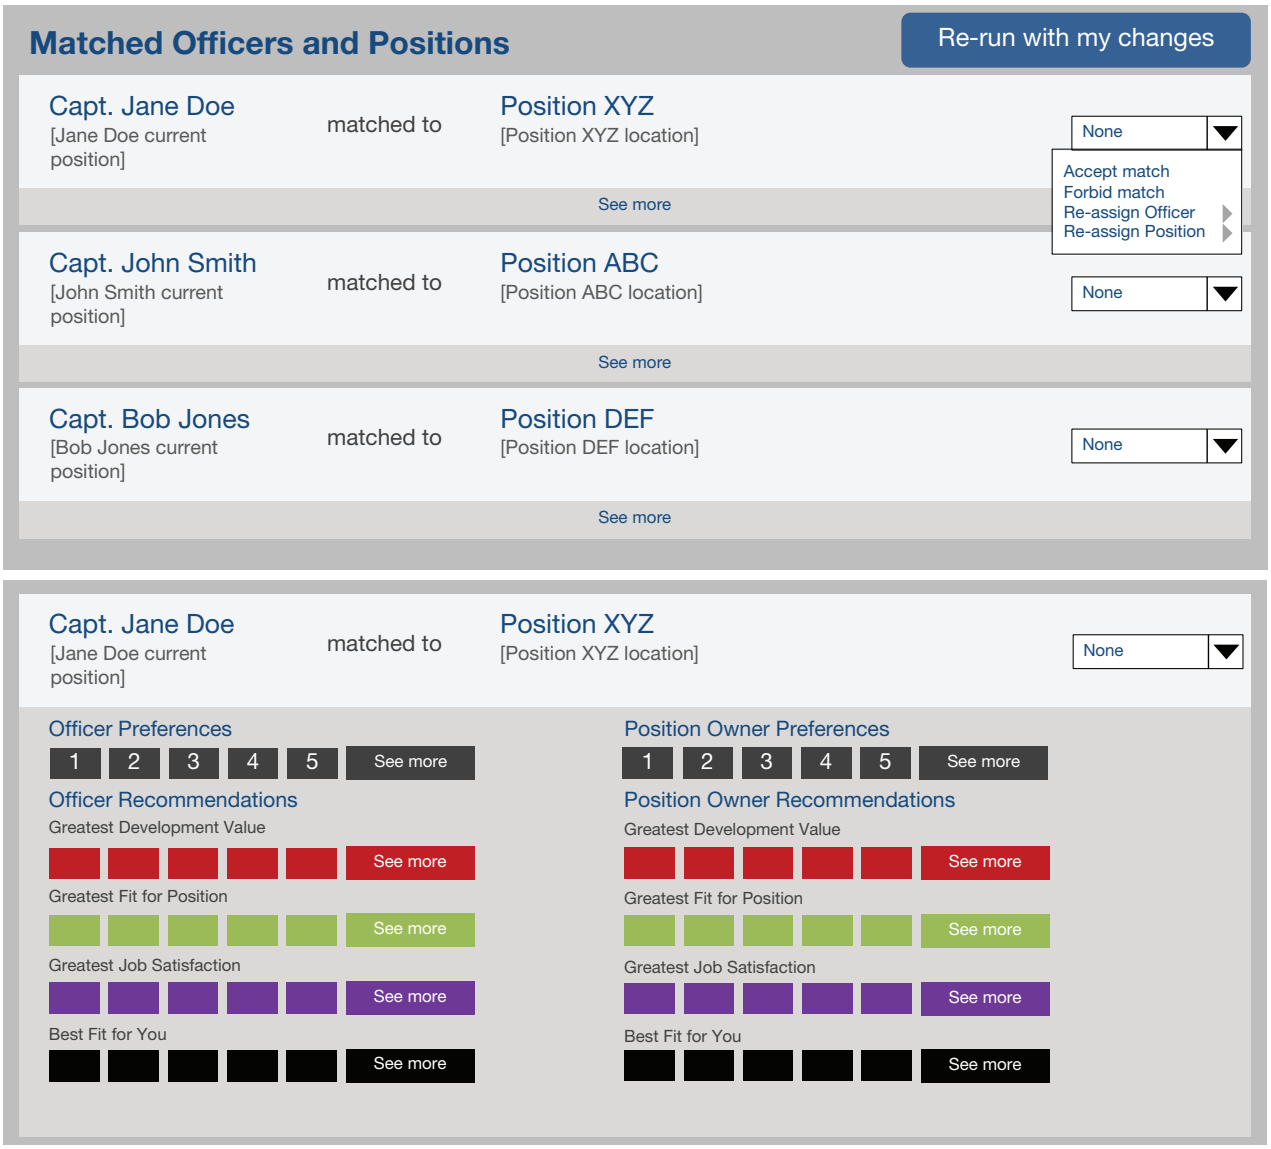

preferences for one or both sides. For example, they could choose to optimize for development value on the officer's side and optimize for performance on the position owner's side.

Discussion and Concluding Remarks

This high-level design analysis of the viability and utility of recommendation systems to enable Talent Marketplace assignment systems supports four key findings. First, with the advent of the Talent Marketplace, the USAF is well positioned to implement ML-enhanced assignment recommendations. The officer and position information needed to enable recommendations are contained in administrative databases and can be augmented by applying NLP techniques to narrative descriptions of duty responsibilities and officer performance contained in OPRs. Besides providing a platform to deliver recommendations, the Talent Marketplace can be used to gather additional information about officer qualifications, position details, and historic preferences. Such addi-

Recommendation engines like those described in this volume use ML to augment, rather than automate, human decisionmaking.

tional information could improve the usefulness of future recommendations.

Second, the themes of officer development, performance, and job satisfaction are a useful and important way to organize the streams of recommendations, because they capture the factors most likely to be important to officers, position owners, and the career field assignment team. Data sources already exist to enable development-based and potentially position-fit–based recommendations, and therefore, the USAF can begin to develop prototype systems to deliver recommendations based on these two themes. However, the USAF currently lacks data to recommend positions based on satisfaction. Talent management channels have likely amassed a list of key dimensions that relate to job satisfaction as part of the USAF's biannual Career Decisions Survey and periodic exit surveys. The next step would be to identify the top job characteristics that relate to satisfaction and begin collecting a set of ratings from current position holders that feed into the Talent Marketplace system. Beyond the value of this information in generating informative recommendations, it would be a key source of feedback to position owners on elements of the work environment that they could improve upon.

Third, recommendation engines like those described in this volume use ML to augment, rather than automate, human decisionmaking. Throughout this series, we have emphasized that ML should function as a decision aid rather than making the decision itself. Furthermore, the ML system may not have complete visibility into the various constraints faced by human decisionmakers and, therefore, should not provide the last word on a decision. In this volume, we have identified many points where humans provided key inputs to inform ML recommendations—for example, allowing officers and position owners to specify the relative importance of development, performance, and satisfaction for their decisions; allowing officers and supervisors to identify development goals that the ML system tries to meet; and allowing assignment teams to apply knowledge of further constraints, not represented by the model, when finalizing assignment decisions. The virtue of a recommendation system is that it allows officers, position owners, and assignment teams to flexibly apply additional knowledge beyond the scope of the model.

Fourth and finally, for most officers, position owners, and assignment teams, the number of options in a single cycle might not be large enough to necessitate personalized recommendations. However, senior HRM decisionmakers increasingly recognize the costs of frequently rotating members to new locations, and they are seeking to transition from a system that matches officers who have been preselected for rotation to a system that selects officers for rotation to facilitate developmental needs (Department of the Navy, 2021, p. 11). Recommendation systems for all entities in a talent marketplace system are an enabling technology to facilitate this increase in the scope of assignment possibilities. Instead of considering only a subset of officers and vacancies each cycle, the Talent Marketplace could consider the projected vacancies over a one- or two-year time horizon, creating significantly more options for matches that facilitate better development and meet individual needs. As the number of potential options increases, so does the need for high-quality recommendations. Thus, the services should begin testing the concepts discussed in this report and exploring enriched data inputs for the next generation of talent marketplace systems.

One important consideration for the USAF to keep in mind is drawn from earlier volumes in this series (see Table 3). In particular, the USAF must carefully consider risks related to accuracy, fairness,

TABLE 3
Outline of Report Series

Volume Number	Report Title	Report Purpose
1	*Leveraging Machine Learning to Improve Human Resource Management: Volume 1, Key Findings and Recommendations for Policymakers* (Schulker, Walsh, et al., 2024)	Overview for senior leaders
2	*Machine Learning in Air Force Human Resource Management: Volume 2, A Framework for Vetting Use Cases with Example Applications* (Walsh et al., 2024)	Framework for how to prioritize ML projects
3	*The Personnel Records Scoring System: Volume 3, A Methodology for Designing Tools to Support Air Force Human Resources Decisionmaking* (Schulker, Williams, et al., 2024)	Technical report on scoring officer records
4	*Safe Use of Machine Learning for Air Force Human Resource Management: Volume 4, Evaluation Framework and Use Cases* (Snoke et al., 2024)	Case study approach to ensure safety of ML systems
5	*Machine Learning–Enabled Recommendations for the Air Force Officer Assignment System: Volume 5,* (Calkins et al., 2024)	ML system to inform officer assignments

NOTE: Current report is highlighted.

and explainability when considering how to use ML to enhance assignment decisions. As described in Volume 4 of this series (Snoke et al., 2024), these dimensions constitute an "iron triangle" that underlies the safe use of ML in HRM. With respect to accuracy, although ultimate decision authority lies with humans, past research has demonstrated the strong propensity for *automation bias*—that is, for people to favor suggestions from decision-support systems even in the face of contradictory evidence (Parasuraman and Manzey, 2010). For this reason, it is important that the recommendation system have high accuracy. Accuracy could be assessed in terms of concordance with expert human decisions. Alternatively, if the Talent Marketplace were instrumented to record measures of satisfaction and job performance, the accuracy of ML systems could be assessed with respect to the outcomes they seek to maximize.

With respect to fairness, the USAF must ensure that the system does not generate assignment recommendations that drive disparate outcomes based on race or ethnicity, gender, and other protected characteristics. This could occur, for example, if the system were trained using historical data with assignment sequences that broke out along demographic lines either directly or indirectly because of moderating variables like Air Force Specialty Code (Lim et al., 2014). A related concern is that because most officers are white males, recommendations may be more accurate for those individuals than for individuals from other demographic subgroups. Finally, with respect to explainability, presenting recommendations in an interpretable manner may allow human decisionmakers to intervene when the model's outputs are incorrect and to detect whether a model violates fairness norms (Adadi and Berrada, 2018). Should the recommendations fail on any of these dimensions, it is possible that they could *increase* the workload of users—as users try to understand and eventually disregard the recommendations—rather than serving as a useful tool to augment decisionmaking.

A final concern is for individuals' privacy. Data contained in administrative databases, OPRs, and other sources, though potentially useful for predicting such outcomes as job performance, are potentially sensitive. Thus, the USAF would need to take care in considering whether to make these data available to position owners in the Talent Marketplace.

Given the findings and considerations presented above, **the USAF should experiment with delivering position and candidate recommendations using recommendation engines.** This may enhance the current Talent Marketplace, and it may reduce risks

associated with transitioning to a true, decentralized marketplace. We suggest (a) the beginning of prototyping for a recommendation system that includes the development and performance themes, as well as a global solution for the career field assignment team and (b) the beginning of data collection that will allow the addition of a job satisfaction theme and an improved position-fit theme. For each prototype, we suggest pursuing a hybrid approach that layers the content-based and collaborative-filtering (and, if applicable, the knowledge-based) approaches together, because it is likely to be more robust than any system on its own. When the prototype is complete, testing can begin within the Talent Marketplace, where recommendations should be evaluated based on their usefulness to officers, position owners, and the assignment team. As mentioned previously, the system is not a replacement for careful decisionmaking by officers, position owners, or career field assignment teams but instead is a tool to help their decisionmaking. The prototypes should therefore be deployed alongside careful messaging emphasizing that they are meant to be decision-support tools only. As prototyping continues, **the USAF should gather new sources of data, such as officers' satisfaction with positions and position owners' satisfaction with officer performance.** Aside from directly benefiting the future performance of recommendation engines, this information can provide valuable information to other USAF human resource management processes.

Finally, **the USAF should experiment with using a matching algorithm to combine officer and position owner preferences.** As implemented, the Talent Marketplace makes information visible and increases the transparency of assignments. Using a matching algorithm is the next step toward transitioning to a true, decentralized marketplace that matches officers to positions based on mutual preference—an approach that has been shown to be beneficial in other markets. However, the matching algorithm must consider the need to meet warfighting requirements and to accommodate special circumstances, such as by-name requests and family medical needs, and the assignment team should be transparent about the process they use and the constraints they face should the USAF shift to an algorithmic approach to assignments.

Although we primarily focused on the utility and feasibility of recommendation engines for USAF officer assignments, we expect that this approach would also be of interest to the other military services. This approach may even be useful beyond the Department of Defense and in civilian markets that use information on preferences to match workers to jobs, such as the National Medical Residency Match.

Notes

[1] This theme may be of high interest to officers because it provides information not available elsewhere except through informal channels. It may also be of high interest to assignment teams for identifying officers who would be least dissatisfied in hard-to-fill positions.

[2] The ADP was a planning form that officers filled out as part of the assignment system before the Talent Marketplace. It allowed officers to input the types of future assignments they desired for their development.

[3] Aside from enabling recommendations, this feedback could provide useful information to assignment teams—for example, by shedding light on why officers attempt to avoid certain assignments and by identifying climate issues within a particular unit or command.

References

Adadi, Amina, and Mohammad Berrada, "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," *IEEE Access*, Vol., 6, 2018. As of December 26, 2023: https://ieeexplore.ieee.org/document/8466590

DAF—*See* Department of the Air Force.

Department of the Air Force, *Total Force Assignments*, Manual 36-2110, August 2, 2021. As of December 26, 2023: https://static.e-publishing.af.mil/production/1/af_a1/publication/dafi36-2110/dafi36-2110.pdf

Department of the Navy, U.S. Marine Corps, "Talent Management 2030," webpage, November 2021. As of December 26, 2023: https://www.hqmc.marines.mil/Portals/142/Users/183/35/4535/Talent%20Management%202030_November%202021.pdf?ver=E88HXGUdUQoiB-edNPKOaA%3d%3d

Gale, David, and L. S. Shapley, "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, Vol. 69, No. 1, 1962. As of December 26, 2023: https://www.jstor.org/stable/2312726?origin=JSTOR-pdf&seq=1

Kendal, Simon L., and M. Creen, *An Introduction to Knowledge Engineering*, Springer London, 2007.

Lim, Nelson, Louis T. Mariano, Amy G. Cox, David Schulker, and Lawrence M. Hanser, *Improving Demographic Diversity in the U.S. Air Force Officer Corps*, RAND Corporation, RR-495-AF, 2014. As of December 26, 2023: https://www.rand.org/pubs/research_reports/RR495.html

McFadden, D., "Conditional Logit Analysis of Qualitative Choice Behavior," in P. Zarembka, ed., *Frontiers in Econometrics*, Academic Press, 1973.

National Academies of Sciences, Engineering, and Medicine, *Strengthening U.S. Air Force Human Capital Management: A Flight Plan for 2020–2030*, The National Academies Press, 2021.

Al-Otaibi, Shaha T., and Ykhlef Mourad, "A Survey of Job Recommender Systems," *International Journal of Physical Sciences*, Vol. 7, No. 29, 2012.

Parasuraman, Raja, and Dietrich H. Manzey, "Complacency and Bias in Human Use of Automation: An Attentional Integration," *Human Factors*, Vol. 52, No. 3, June 2010.

Roth, Alvin E., and Elliott Peranson, "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *American Economic Review*, Vol. 89, No. 4, September 1999. As of December 26, 2023: http://stanford.edu/~alroth/papers/rothperansonaer.PDF

Roth, Alvin E., and Marilda Sotomayor, "Two-Sided Matching," in Robert Aumann and Sergiu Hart, eds., *Handbook of Game Theory with Economic Applications*, Vol. 1, Elsevier, 1992. As of December 26, 2023: https://www.sciencedirect.com/handbook/handbook-of-game-theory-with-economic-applications/vol/1/suppl/C

Schulker, David, Lisa M. Harrington, Matthew Walsh, Sandra Kay Evans, Irineo Cabreros, Dana Udwin, Anthony Lawrence, Christopher E. Maerzluft, and Claude Messan Setodji, *Developing an Air Force Retention Early Warning System*: *Concept and Initial Prototype*, RAND Corporation, RR-A545-1, 2021. As of December 26, 2023: https://www.rand.org/pubs/research_reports/RRA545-1.html

Schulker, David, Nelson Lim, Luke J. Matthews, Geoffrey E. Grimm, Anthony Lawrence, and Perry Shameem Firoz, *Can Artificial Intelligence Help Improve Air Force Talent Management? An Exploratory Application*, RAND Corporation, RR-A812-1, 2021. As of December 26, 2023: https://www.rand.org/pubs/research_reports/RRA812-1.html

Schulker, David, Matthew Walsh, Avery Calkins, Monique Graham, Cheryl K. Montemayor, Albert A. Robbert, Sean Robson, Claude Messan Setodji, Joshua Snoke, Joshua Williams, and Li Ang Zhang, *Options for Applying Artificial Intelligence to Human Resource Management*: Vol. 1, *Substantive Change is Within Reach*, RAND Corporation, RR-A1745-1, 2024.

Schulker, David, Joshua Williams, Cheryl K. Montemayor, Li Ang Zhang, and Matthew Walsh, *The Personnel Records Scoring System*: Vol. 3, *A Methodology for Designing Tools to Support Air Force Human Resources Decisionmaking*, RAND Corporation, RR-A1745-3, 2024.

Snoke, Joshua, Matthew Walsh, Joshua Williams, and David Schulker, *Safe Use of Machine Learning for Air Force Human Resource Management*: Vol. 4, *Evaluation Framework and Use Cases*, RAND Corporation, RR-A1745-4, 2024.

Walsh, Matthew, Sean Robson, Albert A. Robbert, and David Schulker, *Machine Learning and Artificial Intelligence in Air Force Human Resource Management*: Vol. 2, *A Framework for Vetting Use Cases with Example Applications*, RAND Corporation, RR-A1745-2, 2024.

Acknowledgments

We thank Gregory Parsons, Col. Laura King (AF/A1H), and Doug Boerman (AF/A1X) for their input and support throughout the project. We also thank the many RAND colleagues who helped with this work: principally, but not exclusively, Melissa Baumann, Benjamin Gibson, Lisa Harrington, Ignacio Lara, Nelson Lim, Miriam Matthews, Al Robbert, Sean Robson, and Joshua Snoke. Finally, we would like to thank our peer reviewers, Brian Kelly, Caitlin Lee, and Pete Schirmer, whose thoughtful and constructive feedback greatly improved this document.

About This Report

The Department of the Air Force (DAF) has begun to develop and field artificial intelligence and machine learning (ML) systems for myriad mission areas and support functions, including human resource management. ML systems have the potential to accelerate existing decision processes and to enhance decision quality by leveraging data. Further, by allowing the DAF to make decisions at greater speed and scale, ML systems may enable entirely new decision processes.

This research was commissioned by the Director of Plans and Integration, Deputy Chief of Staff for Manpower and Personnel, Headquarters U.S. Air Force and conducted within the Workforce, Development, and Health Program of RAND Project AIR FORCE as part of a fiscal year 2022 project, "Machine Learning Decision-Support Tools for Talent Management Processes." This report is the fifth in a series of five reports originating from this project. The companion reports are (1) *Leveraging Machine Learning to Improve Human Resource Management: Volume 1, Key Findings and Recommendations for Policymakers* (Schulker, Walsh, et al., 2024); (2) *Machine Learning in Air Force Human Resource Management: Volume 2, A Framework for Vetting Use Cases with Example Applications* (Walsh et al., 2024); (3) *The Personnel Records Scoring System: Volume 3, A Methodology for Designing Tools to Support Air Force Human Resources Decisionmaking* (Schulker, Williams, et al., 2024); (4) *Safe Use of Machine Learning for Air Force Human Resource Management: Volume 4, Evaluation Framework and Use Cases* (Snoke et al., 2024). These closely related volumes share some material, including some definitions and descriptions.

RAND Project AIR FORCE

RAND Project AIR FORCE (PAF), a division of the RAND Corporation, is the Department of the Air Force's (DAF's) federally funded research and development center for studies and analyses, supporting both the United States Air Force and the United States Space Force. PAF provides the DAF with independent analyses of policy alternatives affecting the development, employment, combat readiness, and support of current and future air, space, and cyber forces. Research is conducted in four programs: Strategy and Doctrine; Force Modernization and Employment; Resource Management; and Workforce, Development, and Health. The research reported here was prepared under contract FA7014-22-D-0001.

Additional information about PAF is available on our website: www.rand.org/paf/

This report documents work originally shared with the DAF on September 13, 2022. The draft report, dated September 2022, was reviewed by formal peer reviewers and DAF subject-matter experts.

www.rand.org

RR-A1745-5