MATCHING LAB

TARGETED MATCHING STRATEGY OPTIMIZATION
DRIVEN BY DATA AND DOMAIN KNOWLEDGE

WHITE PAPER

Introduction

When matching jobseekers and vacancies in real-world public employment services, most jobseekers tend to be sufficiently served by a limited number of matching strategies that generalize well across the population. Nevertheless, the specific labor market challenges faced by a smaller part of the population typically require more tailored matching strategies. Since it is usually these exact jobseekers that benefit the most from the involvement of public employment services, it is imperative to adequately address their challenges. Therefore, with our Matching Lab, we aim to identify these jobseekers, characterize their labor market challenges, and address these challenges by means of targeted matching strategies that help maximize the number of high-quality matches.

Matching Strategy Optimization

A more tailored matching strategy can be required in various situations. For example, some jobseekers may have a desired occupation that only has a very limited number of vacancies. Other jobseekers may not have the education required by vacancies for their desired occupation, despite having enough experience to compensate for that. Such cases require matching strategies that, for instance, allow for alternative, similar occupations to be accounted for, emphasize skills or experience over education, or increase the typical radius of the physical (travel) distance between jobseekers and vacancies.

In practice, it is not immediately obvious exactly which jobseekers require a more tailored approach, which job market challenges should be addressed by such an approach, and how such an approach should address these challenges. The vast amounts of data typically involved in employment matching hold valuable clues that can help answer such questions. Furthermore, domain knowledge of experts can help to put such insights into the right context. Therefore, we have devised a seven-step process to help optimize matching strategies through the combination of data-driven insights and domain knowledge by involving both data scientists and domain experts. A schematic overview of this process can be found in Figure 1.

Figure 1: An overview of the Matching Lab process.

MATCHING LAB
TARGETED MATCHING STRATEGY OPTIMIZATION

Step 1: Observe

The first step in our matching strategy optimization process is to observe the labor market in order to understand the historical and current trends and patterns in the labor market. To this end, we clean, standardize, and enrich the available labor market data and load it into our labor market dashboards. These dashboards visualize the relevant data in informative graphs that allow for further exploration and dissection of the data.

VDAB, the Belgian public employment service for the region of Flanders, has recently used Matching Lab in order to optimize their matching strategies. Figure 2 shows an example of a dashboard that provides insight into some anonymized and randomly sampled data for the Flemish labor market.

Apart from exploring general trends and patterns in the labor market, we evaluate the number of matching vacancies over time for each jobseeker through an applicable matching strategy. In case of an existing employment matching setup, this matching strategy is typically the comparably generic strategy that is currently applicable to a jobseeker. In case no employment matching setup is currently in place, we configure a matching setup with a generic matching strategy that is based on our best practices, insights obtained from the available data, and domain knowledge of labor market experts.

Step 2: Filter

Once the labor market situation has been assessed, we proceed to narrow down to the jobseekers who are viable candidates for more tailored matching strategies. These jobseekers are the ones who are systematically not served well by the general matching strategy.

Here, we typically set out to identify the jobseekers who tend to get too few matching vacancies. Other interesting application

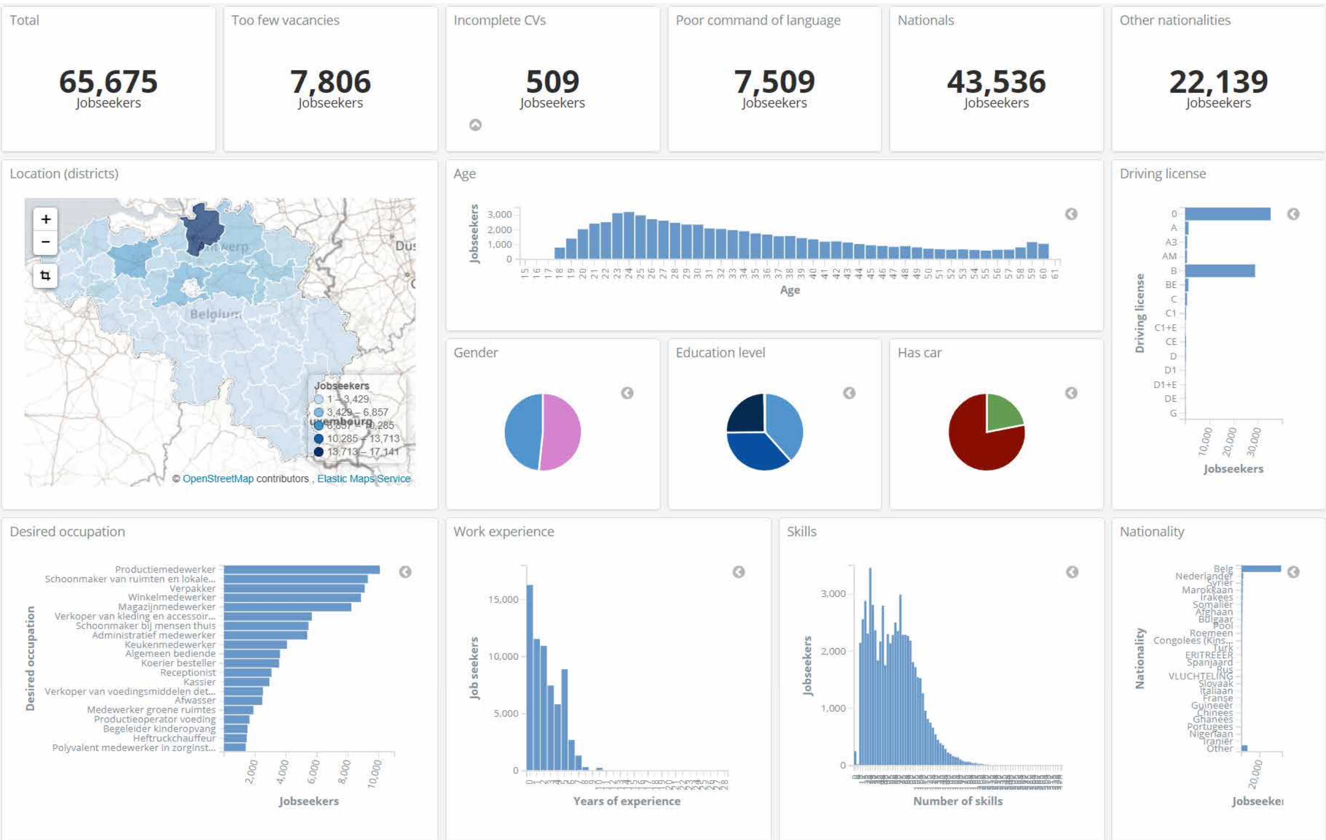

Figure 2: An example of a dashboard that provides insight into labor market data through aggregate statistics, interactive maps, and charts that visualize distributions of values for key properties.

scenarios for more tailored matching strategies include, for example, addressing the needs of jobseekers who tend to get too many (possibly less relevant) matching vacancies, or of jobseekers whose matching vacancy counts fluctuate heavily over time.

In order to identify the jobseekers who get too few matching vacancies over time, we typically periodically observe the number of matching vacancies for each jobseeker. Then, we focus on those jobseekers having a systematically low number of matching vacancies. A typical example of such a selection criterion is to select those jobseekers that have at most 10 matching vacancies, for at least 75% of all observations. An alternative approach here would be to simply focus on the jobseekers who, over the observed period of time, have an average number of matching vacancies below a certain threshold.

Step 3: Identify Groups

The selected jobseekers all have one thing in common, i.e., the fact that they are not optimally served by the general matching strategy. Nevertheless, this set of jobseekers is typically very diverse and is consequently associated with a myriad of labor market challenges that contribute to the limited applicability of a more generic matching strategy. Each of these challenges may require a different, targeted matching strategy. In order to be able to address these challenges as effectively and efficiently as possible, it is recommended to identify homogeneous groups of jobseekers who have similar challenges on the labor market.

In practice, groups of jobseekers who share similar characteristics tend to face similar challenges on the labor market. Therefore, we set out to group our considered jobseekers based on their observable characteristics, with a specific focus on the relevant dimensions of age, gender, nationality, location, mobility, education level, competences, languages, (desired) occupations, and work experience. We aim to segment the jobseekers in such a way that, from a data perspective, the jobseekers within each group are as similar to one another as possible, whereas the groups are as dissimilar to one another as possible. A machine-learning technique that is particularly suited for this task is clustering, which aims to identify groups in high-dimensional data while maximizing within-group similarities and minimizing between-group similarities.

One of the considerations when clustering the jobseekers who are not served well by the general matching strategy is how to best represent the high-dimensional data of these jobseekers. Here, we typically experiment with various representations of non-numeric values, which include hashing or vectorizing these values using, e.g., one-hot encoding. Moreover, we consider reducing the dimensionality of the data by applying techniques like singular value decomposition. Such dimensionality reduction techniques transform the data into a lower-dimensional space where each dimension is a (typically linear) combination of the original dimensions, explaining a distinct and as large as possible part of the variance in the data.

Another consideration in the clustering process is which clustering method to use. We typically experiment with various algorithms, including k-means clustering, agglomerative hierarchical clustering, the neural gas method (inspired by self-organizing maps), CLARANS (medoid-based clustering), and BIRCH (large-scale hierarchical clustering). Additionally, we experiment with various parameterizations for these algorithms.

We thus produce various sets of clusters by combining our considered data representations, clustering algorithms, and their parameterizations. With the help of domain experts, we evaluate these sets of clusters in terms of to what extent they group jobseekers with similar challenges on the labor market in a useful and actionable way. Based on these evaluations, we then proceed to work with the most useful set of clusters, representing distinct groups of jobseekers who require more tailored matching strategies.

For VDAB, the most useful groups of jobseekers were produced by k-means clustering on hashed data.

Step 4: Define Tuning Cases

The identified groups of jobseekers who are not optimally served by the general matching strategy form the starting point for the definition of tuning cases. A typical tuning case characterizes a group of jobseekers, their specific labor market challenges, and potential ways of addressing those challenges. In the characterization of the jobseekers covered by a tuning case, we pay special attention to what distinguishes these jobseekers from other jobseekers – e.g., a very distinct distribution of age, gender, education level, field, and/or geographical location. Furthermore, we complement this characterization with a persona, i.e., a fictive jobseeker that represents the typical jobseeker for this tuning case.

Even though the groups of jobseekers identified in step 3 are clearly distinct from a data point-of-view, their specific challenges on the labor market may overlap and/or require similar matching strategy tweaks. Therefore, we work with domain experts and caseworkers in order to consolidate the identified groups of jobseekers into actual tuning cases, by leveraging domain knowledge to interpret the data-driven findings.

An example of a specific, but very actionable tuning case that has emerged from this process for VDAB is the tuning case of highly-educated young men looking for a job in the field of graphic design in the districts of Ghent, Leuven, and Antwerp. The limited availability of vacancies in this field requires alternative matching strategies that, for instance, broaden these jobseekers' horizon in terms of occupations that are closely related in terms of required skills, education, or career-switching opportunities.

Step 5: Fine-tune Matching Strategy

For each tuning case, we optimize the matching strategy by means of an iterative and interactive process. In close collaboration with domain experts, we optimize the matching strategy for a handful of actual jobseekers who form an accurate and balanced representation of the tuning case. These representative jobseekers can be either hand-picked or algorithmically

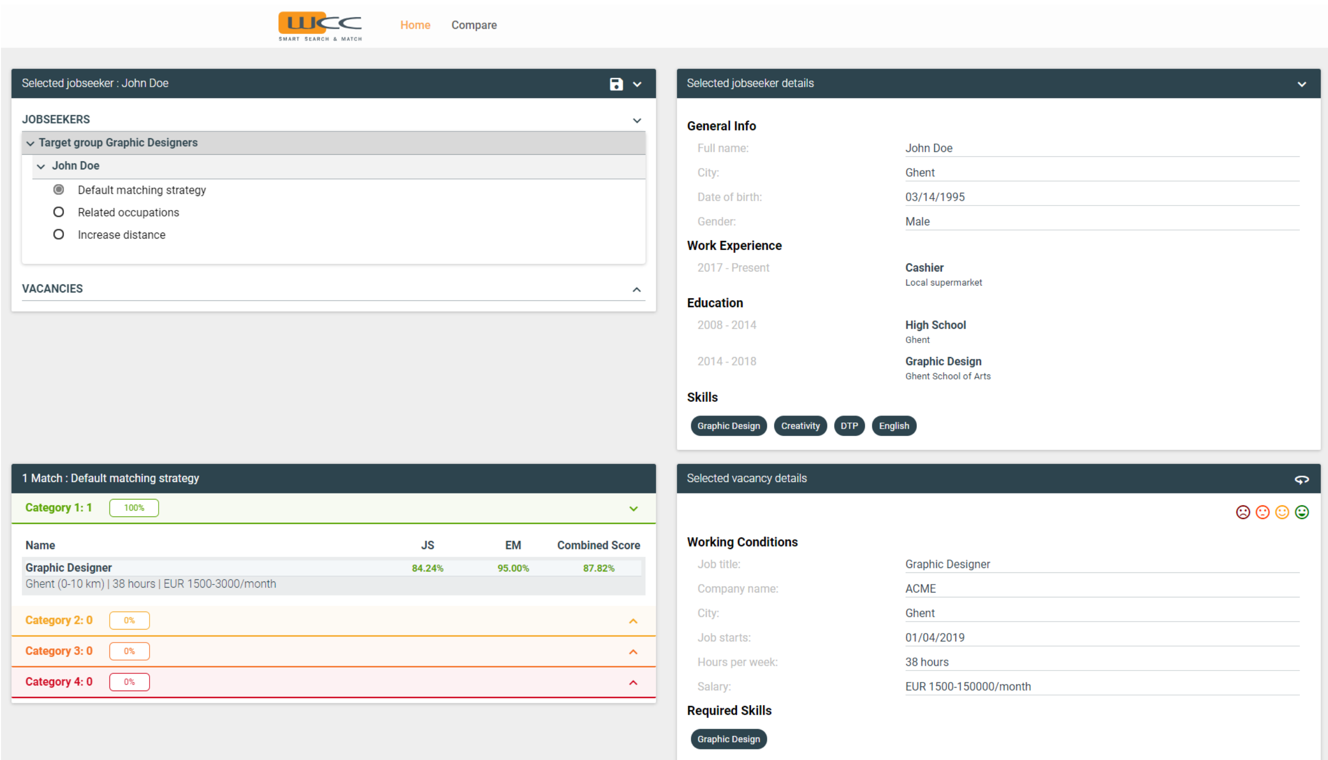

Figure 3: Our tuning tool that provides qualitative insights into how a selection of jobseekers is affected by fine-tuning of the Employment Platform. For each considered jobseeker, the tool allows for an inspection of their relevant properties, as well as for an inspection and evaluation of their matching vacancies as returned by means of different matching strategies.

selected based on equally-spaced percentiles of the distance of jobseekers to the center of the group, i.e., of the distances to the "average" jobseeker in the group of jobseekers covered by this tuning case.

When working on a specific tuning case, we investigate why the selected representative jobseekers are not served well by the general matching strategy, and how we can improve this by means of a more tailored strategy. When employment matching is performed using our Employment Platform, we execute this fine-tuning process by first creating a target group that targets all jobseekers covered by the tuning case. Then, we define one or more perspectives that specify the more tailored matching strategies that address the target group's specific matching challenges. We provide insight into the effects of these perspectives on the representative jobseekers by means of a tuning tool (see Figure 3), designed specifically for this purpose.

Step 6: Evaluate Qualitatively

Once the matching strategy for a specific tuning case has been fine-tuned to address the targeted jobseekers' specific challenges, we evaluate the merits of this matching strategy compared to the more general matching strategy in a qualitative way. We do this by evaluating the individual vacancies returned by the fine-tuned matching strategy for the selected handful of representative jobseekers.

In our qualitative evaluation, we pay special attention to whether or not the tailored matching strategy returns the vacancies we expect it to return, and whether or not it does not return the vacancies we do not expect it to return. Moreover, we evaluate the unexpected match results (if any), in terms of whether or not they make sense. If this analysis indicates that there still is room for improvement, another cycle of tuning (step 5) and evaluation (step 6) may be required. If the qualitative analysis has a positive outcome, we can proceed with quantitative evaluation.

Step 7: Evaluate Quantitatively

In order to quantify the overall effects of using a tailored rather than a generic matching strategy for a specific group of jobseekers targeted by a tuning case, we characterize the (dis)similarities between the match results for all jobseekers in this group, as obtained by applying these two strategies. We do this for individual jobseekers, as well as across all considered jobseekers.

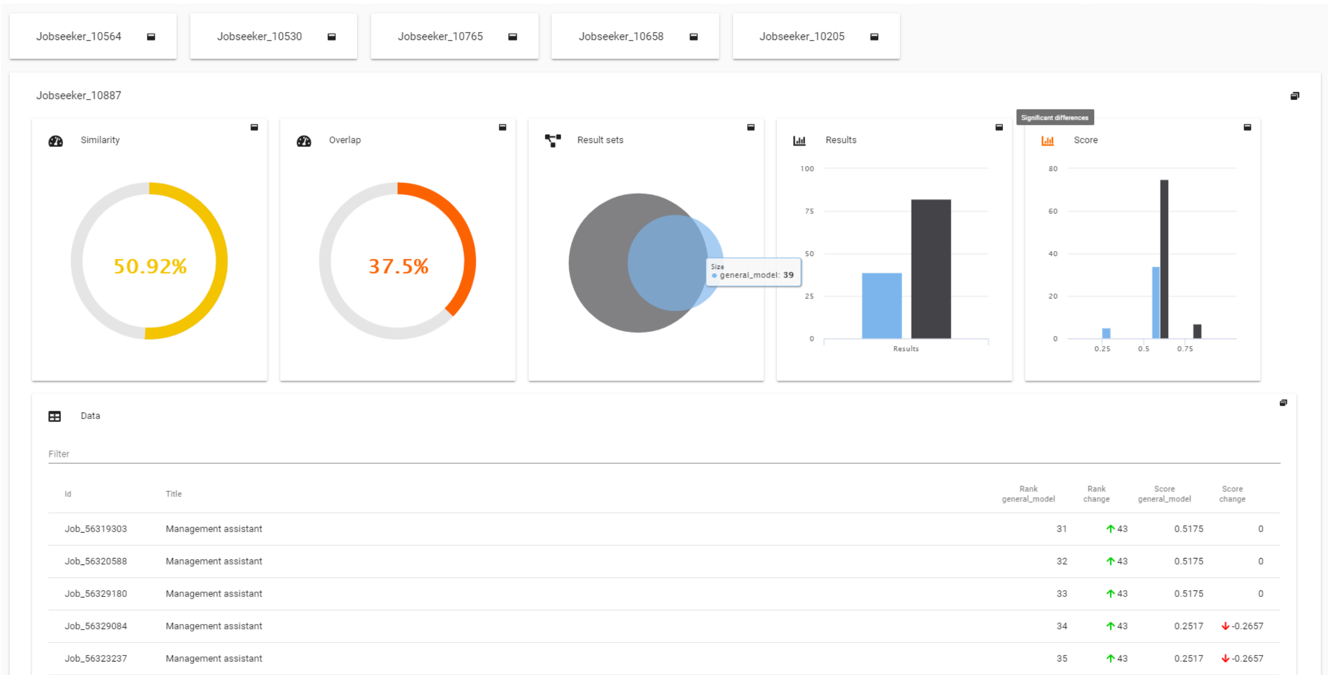

Figure 4: An example of a part of a quantitative evaluation of the vacancies returned by two matching strategies for a single jobseeker.

For each individual jobseeker, we use various metrics to compare the two lists of vacancies that stem from the two respective matching strategies when performing a match for the jobseeker. Then, we aggregate these evaluation metrics over all jobseekers in the group in order to allow for an evaluation of (dis)similarities between match strategies for the group as a whole. If the overall statistics thus obtained show that the more tailored matching strategy has the intended effect, the tuning can be considered complete. If not, further cycles of tuning (step 5) and evaluation (steps 6 and 7) may be required.

Comparing Match Results for a Single Jobseeker

We consider various metrics when comparing the vacancies for an individual jobseeker, as returned by two different matching strategies. In our comparison, an example of which is depicted in Figure 4, we analyze the similarity, overlap, number of results, and scores of the vacancies in the lists of match results generated by the considered matching strategies.

Similarity

The similarity s between two lists of match results i and j is computed as the harmonic mean of the similarity s_{ij} of list i to list j and the similarity s_{ji} of list j to list i, i.e.,

$$s = 2 \cdot \frac{s_{ij} \cdot s_{ji}}{s_{ij} + s_{ji}} .$$

The similarity s_{ij} of list i to j is a ranking-based similarity measure, i.e., the normalized discounted cumulative gain of list i in terms of list j. This similarity measure takes list j as a ground truth and then computes the extent to which the order of items t_i in list i corresponds to their associated relevance $r_{t_i j}$ in list j. In order to emphasize similarities in the higher regions of a ranking, the relevance $r_{t_i j}$ of item t_i in list j is reduced logarithmically proportionally to its position $p_{t_i i}$ in list i, i.e.,

$$s_{ij} = \frac{\sum_{t_i \in i} \frac{2^{r_{t_i j}} - 1}{\log_2\left(p_{t_i i} + 1\right)}}{\sum_{t_j \in j} \frac{2^{r_{t_j j}} - 1}{\log_2\left(p_{t_j j} + 1\right)}}, \quad \begin{array}{l} 0 \le r_{t_i j} \le 1, \quad 0 \le r_{t_j j} \le 1, \\ 1 \le p_{t_i j} \le |\mathrm{i}|, \quad 1 \le p_{t_j j} \le |\mathrm{i}| . \end{array}$$

The relevance $r_{t_i j}$ of item t_i in list j is 0 for all items not in list j, and is the harmonic mean of its position $p_{t_i j}$ (scaled between 0 for the lowest position and 1 for the highest position in the list) and match score $m_{t_i j}$ in list j for all other items, i.e.,

$$r_{t_i j} = 2 \cdot \frac{\frac{|\mathrm{j}| - p_{t_i j} + 1}{|\mathrm{j}|} \cdot m_{t_i j}}{\frac{|\mathrm{j}| - p_{t_i j} + 1}{|\mathrm{j}|} + m_{t_i j}}, \quad 1 \le p_{t_i j} \le |\mathrm{j}|, \quad 0 \le m_{t_i j} \le 1, \quad \forall t_i \in j .$$

Overlap

The extent to which lists i and j overlap is computed as the number of items that occur in both lists, expressed as a fraction of the total number of unique items in the combined lists. In other words, the overlap o_{ij} of lists i and j is defined as the size of the intersection of lists i and j, proportional to the size of the union of lists i and j, i.e.,

$$o_{ij} = \frac{|\mathrm{i} \cap \mathrm{j}|}{|\mathrm{i} \cup \mathrm{j}|} .$$

Results

In order to allow for a comparison between lists i and j in terms of their results, we evaluate the number of results in both lists, i.e., $|i|$ and $|j|$, respectively. Furthermore, for each item that occurs both in list i and in list j, we compute the change in rank and score from list i to list j. Last, we assess the statistical significance of these changes in ranks and scores from list i to list j by means of a paired, two-sided t-test with a significance level of 0.05.

Scores

We use a set of descriptive statistics in order to characterize the distribution of match scores for the items in each list. These descriptive statistics include the mean, the standard deviation, the first quartile, the median, the third quartile, the scores at various quantiles, and the frequencies of scores in various ranges.

Comparing Match Results across Jobseekers

We characterize the overall similarity, overlap, number of results, and the scores of match results returned by two match strategies across all considered jobseekers by means of descriptive statistics. Furthermore, we assess the statistical significance of the changes in ranks and scores across all jobseekers by means of a paired, two-sided t-test with a significance level of 0.05. An example of such an overall comparison can be found in Figure 5.

Even though the overall descriptive statistics are arguably the most useful ones for assessing the (dis)similarities between the results returned by two matching strategies at a glance, the (dis)similarities for some individual jobseekers may warrant some closer inspection as well. Therefore, we use the computed overall descriptive statistics in order to detect jobseekers who may require closer inspection, as their similarity score sets them apart from the other jobseekers as a statistical outlier.

In order to detect statistical outliers in a group of jobseekers, we first compute the interquartile range (IQR) of the similarity scores of these jobseekers, i.e., the difference between the first and third quartile of these scores. Each jobseeker with a similarity score between 1.5 and 3 times the IQR below the first quartile or above the third quartile, is marked as an outlier. Each jobseeker with a similarity score that is more than 3 times the IQR below the first quartile or above the third quartile, is marked as an extreme outlier.

Figure 5: An example of a part of a quantitative evaluation of the vacancies returned by two matching strategies across a group of jobseekers.

Beyond Matching Lab

Matching Lab does not necessarily end once all jobseekers that require a more tailored matching strategy are served by one – it is only a first step towards continuous optimization of matching strategies, guided by data and domain knowledge. Labor market conditions may change over time, due to both external factors and the possible (indirect) effects of the optimized matching strategies. Therefore, we recommend a periodic repetition of the Matching Lab cycle. Furthermore, in between Matching Lab cycles, a new jobseeker who is not served well by a general matching strategy can benefit from a more tailored strategy by analytically assigning this jobseeker to one of the previously identified groups, i.e., the group that has the lowest distance from its center to the jobseeker. Last, it is important to note that Matching Lab is not limited to optimizing the matches for jobseekers. A similar process can be applied to vacancies or other matchable entities on the labor market.

We have shown how our Matching Lab leverages both data and domain knowledge in order to help identify jobseekers that are not optimally served by a general matching strategy, characterize their labor market challenges, and address these challenges by means of targeted matching strategies. This can help public employment services provide significantly better services to those jobseekers that need it the most.

The first Matching Lab cycle of VDAB has already shown to significantly reduce the number of Flemish jobseekers with a consistently low number of matching vacancies. By combining data-driven insights with domain knowledge, these specific jobseekers can now be provided with additional vacancies that, e.g., are a bit further away, require a similar skill set, or provide realistic career-switching opportunities.

Nevertheless, Matching Lab is only a first step in extracting practical value from data and domain knowledge through analytical processes. One can tap further into the potential of data by improving matches through the incorporation of machine-learning techniques like vector-based matching in order to capture, for example, a goodness of fit in terms of intangible and latent concepts like company culture, personality traits, or potential interest in vacancies based on clicking behavior.

About WCC

Our vision

People in organizations make decisions. In the markets we focus on, those decisions profoundly impact people's lives. To make the right decisions in an increasingly complex world, it is necessary to have excellent software. That is what drives us at WCC: enabling people to make better decisions.

Our mission & strategy

WCC wants to give people the answers they need, not just the ones they asked for. We thrive on developing software that can connect, combine, and make sense of large amounts of data stored in different systems. Software that can communicate with the users in a human way, and that delivers superior results so our customers can make a difference. We call this "software that matters". But great software alone is not enough to get the best results. What sets WCC apart is the combination of remarkable software with in-depth knowledge of our customers' business. That is why business and implementation consultancy is an important part of our strategy. We focus on two markets: Employment and Identity.

Our products and services

The core of the Employment market is matching people with sustainable jobs effectively and efficiently. WCC has proven to be unequalled in doing just that. Our Employment Platform, which combines unique search and match capability with advanced gap analysis and referral to the right measures, delivers superior strategic value to our customers. Many of the world's largest employment and staffing organizations use our products and expertise, including Randstad, Robert Half, and the public employment services of Germany, France, and the Netherlands.

The security needs of the Identity market are stringent. Border management and law enforcement agencies face the challenge of quickly and accurately identifying people from huge amounts of data spread over many different databases and formats. WCC's software incorporates the necessary evidence-based algorithms, such as multi-cultural name matching, to make correct identifications. HERMES, our API/PNR solution, adheres to industry standards and is easy to implement and operate. Our customers include UNHCR and the European Union.

WCC
Zonnebaan 19
3542 EA Utrecht
The Netherlands
T: +31 (0)30 750 32 00

info@wcc-group.com
www.wcc-group.com

WCC Services US Inc.
228 Hamilton Avenue
Suite 300, Palo Alto
CA 94301, USA
T: +1-888-922 9224