serco

Speeding Up Applications for the Affordable Care Act with AI

Serco's Solutions Team integrated Rosette to affordably and efficiently process applications for coverage under the Affordable Care Act.

Imagine what it takes to collect, process, and adjudicate 11 million applications every year, with the vast majority happening within a surge period of 135 days. That is what Serco does for the Affordable Care Act (ACA, also known as "Obamacare")—under the HHS Centers for Medicare and Medicaid Services Eligibility Support program.

As one of the world's leading service companies, Serco delivers vital support to government customers in the U.S. and Canada, and commercial customers worldwide. This is one example of the team's rapid response and large program support, especially for critical citizen health programs. Serco's knowledge workers ensure that eligible Americans are able to afford the healthcare they need. Each year, for the year-end 45-day enrollment period and the following 90 days, Serco's knowledge workers are on deadline to process and adjudicate the eligibility of millions of applications.

Serco can apply a similar response and result for emergency health response requirements.

The Challenge

To meet this massive challenge, Serco typically ramps up from 700 trained, off-season workers to over 5,000 for the peak 45-day open enrollment period, which extends from November to December. However, getting

completed applications from initial application to final adjudication is not a simple task.

It starts with a person submitting an application at HealthCare.gov or by mail. Suppose they attest to an annual household income of $40,000 per year. If there is more than a 25% variance in income when compared to IRS or Social Security databases, or if other eligibility factors apply, their application will require manual adjudication. (The majority of applications fall into this category.) Although the application will be provisionally accepted, it needs to be adjudicated within 90 days, pending the receipt of supplemental documentation, such as:

- Social Security cards or birth certificates

- Document numbers for eligible immigrants who want health coverage

- Paystubs, W-2 forms, or other information about household income

- Policy information for any current health coverage

- Information about any employer-provided health coverage available to the applicant.

Mailed-in documentation—such as a spouse's prior year 1040 income tax form—is scanned (using optical character recognition, or OCR), and then must be associated with the correct application.

"The largest volume transaction is 'person association,'" Nathaniel Palmer, Chief Architect at Serco, said. "We deal with myriad rules and different potential relationships, such as between applicant and tax filer, that aren't obvious. Sometimes there's no way to uniformly match correspondence to an issue."

Although names on submissions are all written in English, U.S. resident names come from all over the world. Phonetic transcription errors, nicknames, and cultural naming conventions—such as two-word Hispanic family names—stymie humans and traditional search algorithms.

> "The largest volume transaction is 'person association.' We deal with myriad rules and different potential relationships, such as between applicant and tax filer, that aren't obvious. Sometimes there's no way to uniformly match correspondence to an issue."

Overcoming these matching issues through manual review is difficult and subjective, and consumes many worker hours. Serco is paid a fixed price per task in processing each application, no matter how long the task takes to complete. To ensure that this program does not fall into the red, Serco must correctly estimate the time for each task, and find ways to finish applications in less time.

Sustaining a cash-flow positive program, however, is not the sole measure of success for Serco. "We have program integrity issues that we strive to maintain. We have to demonstrate that every consumer was treated the same way, and that no one got short shrift. We apply the same rules to everyone, and the outcomes are

consistent," Palmer said. "Every decision is appealable, so our adjudication has to stand up to third-party appeal, and the process has to stand up to third-party audit."

To meet these stringent requirements, Serco needed a solution that would:

1. Reduce costs by automating more complex, subjective work (like "person association")

2. Maintain program integrity through a fair and equitable process that could stand up to third-party audit.

The Solution

To decrease the amount of manual labor per task, we sought to expand the volume of what can be automated.

"Some tasks are easy for a person, but hard for a computer, like 'Is this form signed or not?'" Palmer said. "But answering, 'Is that the same person or not based on a name?' is subjective, and the amount of time for a human can vary greatly, and accuracy varies greatly.

"That's where being able to apply sophisticated software makes a huge difference. We are constantly moving the needle of automating that which is complex."

> Answering, 'Is that the same person or not based on a name?' is subjective, and the amount of time for a human can vary greatly, and accuracy varies greatly. That's where being able to apply sophisticated software makes a huge difference. We are constantly moving the needle of automating that which is complex."

Palmer and his team evaluated several name matching solutions, including Rosette Name Indexer and the basic Apache Solr open source search engine distribution. The team built proofs of concept, running them through more than 100 test cases to determine their accuracy in performing pairwise name matching. In the analysis of test results, Palmer said, "The difference [in accuracy] was significant enough to make the case to go with Rosette."

Rosette Name Indexer's accuracy comes from knowledge of how names are formed and how they vary in each language. This knowledge is built into a variety of algorithms. Depending on the names, different algorithms are invoked to match names of people, organizations, and places across 15+ languages and 13 types of name variations, such as nicknames, typos, and swapped name components.

For every pair of names, Rosette outputs a match score ranging from 0 (no match) to 1 (perfect match), and is transparent in labeling the algorithms that produced that score. Other fields, such as Social Security number and date of birth, can be custom weighted and factored into the score, too. Each Rosette user sets the threshold, such as 0.8, above which names are considered "a match."

Rosette was first integrated to assist Serco in verifying applicants who had not responded to requests for additional information. Serco knowledge workers place about 10 million outbound phone calls every year, but their success rate in reaching someone is less than 25%. In the event that there is no response to their entreaties in the 90 days after an application has been submitted, they are dubbed "radio silent ever."

> "We have to confirm we really haven't heard from them, so we need to query the name. From the government's standpoint, it wants as many eligible people as possible to get coverage. That's where Rosette comes in now. It enables automation and bridges the gap."

"We have to confirm we really haven't heard from them, so we need to query the name. From the government's standpoint, it wants as many eligible people as possible to get coverage," Palmer said. "That's where Rosette comes in now. It enables automation and bridges the gap. We couldn't apply automation there without Rosette."

The Impact

Following a one-day training and tuning session, Serco integrated Rosette into its production system. Within days of deploying Rosette, Serco noted a substantial improvement in the accuracy of name matches across its workflow, as well as a marked decrease in the time required to perform name matches when processing applications.

> Rosette "has saved in the 100,000s of minutes this period. The reality is [without Rosette], we wouldn't be able to apply automation."

With Rosette automating the verification of "radio silent ever" applications only since Jan. 1, 2020, Palmer estimated in April 2020 that Rosette "has saved in the 100,000s of minutes this period. The reality is [without Rosette], we wouldn't be able to apply automation."