Business Transformation & Operational Excellence World Summit Nomination for:

- Best Achievement in Technology Enabled Process Automation Award

Caesars Entertainment Analytics

Rias Attar, PMP, PMI-ACP, Scrum Master, LSSBB
VP, EPMO and Strategic Initiatives
November 26th 2018

Synopsis:

Caesars Entertainment has multiple machine learning/AI initiatives recently completed as part of a larger project to automate and improve processes. We are highlighting two of these efforts here.

The data science team developed cloud-based machine learning models to improve the customer valuation and marketing campaigns processes. The solution deployed seamlessly integrates with centralized enterprise wide marketing campaigns scoring process across all properties and the entire active customer database. The solution alleviates the need for manual adjustments and resulted in double-digit accuracy gains over the prior valuation.

An automated process was developed for curation of Point of Sale (PoS) item taxonomies using AI-based Text Mining and pattern recognition. Due to complex history of large acquisitions and accompanying variety of PoS backend systems as well as inconsistency in manual entry, there is a constant challenge of Master Data management of PoS items (Retail, Food & Beverage) within Caesars enterprise. Manual cleaning and curation are not scalable; therefore, Caesars' Data Science team streamlined this process via automation and leveraged AI based solutions to perform this daily/weekly across ~500K individual items. The solution developed leverages customer AI, text mining and pattern recognition algorithms to constantly maintain complete taxonomy of PoS items sold across Caesars outlets. Manual efforts invested into adjusting PoS data and interpreting data were reduced by ~80%.

Cloud-based Machine Learning models to improve customer valuation and marketing campaigns process

Current State

Historically Gaming companies have approached customer valuation as a static scoring process (similar to customer lifetime calculation for banking customers). This results in large accuracy variances in what is predicted vs. what is the ground truth (actual) in future worth of a customer.

Inaccurate future prediction can result in inefficient marketing spend (in terms of offers, customer campaigns etc.) and overreliance on manual adjustment.

Caesars Entertainment's Data Science team addressed this fundamental problem by building advanced machine learning models in its cloud environment. The modeling effort leveraged state of the art open sourced machine learning packages along with a custom-built data lake.

Developed Solution Details

The deployed solution seamlessly integrates with centralized enterprise wide marketing campaigns scoring process across all properties encompassing the entire active customer database. Additionally, current solution alleviates prior manual adjustment process, as it is significantly more accurate in terms of customer valuation prediction.

The entire modeling stack and associated data lake is purpose-built in a cloud environment which makes it easy to scale out this process to keep in line with growth experienced in the customer database. Implementation and working in a cloud environment empower a "fail fast" development framework, allowing for a more efficient and streamlined modeling efforts.

The model deployed incorporates several underlying models. After hundreds of iterations (made possible through the "fail fast" environment), the ensemble model encompasses a variety of regression and classification supervised machine learning frameworks including neural networks, boosted trees, and clustering techniques.

Key results

New Machine Learning model resulted in double-digit accuracy gains over the prior valuation and has been extensively test marketed using cloud-based testing frameworks.

Due to substantially more accurate predictions, we have seen a 300 to 900 bps improvement in overall marketing efficiency.

Future outlook

Based on success experienced through this solution for core marketing campaigns, in 2019 the team is now deploying this solution for all investment decisions within marketing campaigns. This impacts a sizable portion of the overall marketing spend for Caesars Entertainment.

Automated curation of Point of Sale (PoS) item taxonomies using AI based Text Mining and pattern recognition

Current State

Due to complex history of large acquisitions and accompanying variety of PoS backend systems as well as inconsistency in manual entry, there is a constant challenge of Master Data management of PoS items (Retail, Food & Beverage) within Caesars enterprise. Multiple efforts have been undertaken in the past to build a manually cleaned & curated taxonomy. But this is not scalable given constant changes in product offerings.

Caesars' Data Science team was tasked to completely overhaul this manual and expensive process and to streamline this process via automation and to leverage AI based solutions to perform this daily/weekly across ~500K individual items. The process needed to be robust in-light of updates and inclusion of additional items frequently.

Developed Solution Details

The solution developed leverages customer AI, text mining and pattern recognition algorithms to constantly maintain complete taxonomy of PoS items sold across Caesars outlets.

AI based algorithms discover latent word associations, remove duplicate entries based on text analysis and maintain an intelligent hierarchical system of relationships within entire universe of PoS items.

The process is completely automated with a custom front-end to integrate with multiple PoS systems and accept individual operator's manual inputs. The system also integrates with Caesars Enterprise Data warehouse to keep an up to date taxonomy and item hierarchy.

This data is then used by multiple teams to perform reporting and financial reconciliation, providing decision makers accurate insights across key decision areas (sales, volume, pricing, inventory etc.)

Key results

The immediate impact of the AI-based models yielded a clean, usable, structured dataset. The process is streamlined and fully automated, allowing for on-demand large-scale analysis, empowering and enabling key stakeholders to improve their efficiencies. The average turnaround timeframe reduced from weeks to days. The department previously responsible for manually curating data can now focus their time on higher value projects. Manual efforts invested into adjusting PoS data and interpreting data were reduced by ~80%.

Future outlook

This solution is now rolled out to multiple geographies (US, Ex-US) and larger end user groups (including front line operations) as a core part of their day-to-day workflow. Any future discussion on new PoS systems are also being evaluated based on ease of integration with internally developed solution, as this results in continued adoption and shorter time frame to deliver value.