#2 Aria

#18 Vdara

#13 Cosmopolitan

#19 Hollywood #22 Westgate

#11 Paris

#6 New York New York

#20 Excalibur

#24 Bally's #21 Tropicana

FOLIESBERGER

#1 MGM

#38 Encore

Learning to Rank

with Eric

#14 Luxor





- 1. Introducing the problem
- 2. Evaluating performance
- 3. Interpreting the model
- 4. Designing the machine
- 5. Deploying to production

Historical attempts have not succeeded

R

Deployment

Points = $w_1 * \text{stars} + w_2 * \text{rating} + w_3 * \text{price} + w_4 * \text{spread} + \dots$

app_id	experiment_name	start_dt	ttm_lift
hotelstorm-www	Config:hs.exp.search.ranking_algo_points	Nov 2017	-0.06992
hotelstorm-www	Config:hs.exp.search.ranking_algo_points	Jan 2018	0.01064
hotelstorm-www	Config:hs.exp.search.ranking_algo_points	Feb 2018	0.0141
ALL-whitelabel-web	Config:testarossa.searchResults.rank.weights	Feb 2018	0.00466
hotelstorm-www	Config:hs.exp.search.ranking_algo_points_v2	Mar 2018	-0.03244
hotelstorm-www	Config:hs.exp.search.ranking_algo_points_v3	Apr 2018	0.01407

Introduction

Evaluation

Interpretation

Machine Learning

Ranking is a complex problem!





Our solution balances all factors





#22 Reflect

#28 Hyatt Ziva

#2 Grand Fiesta Americana Coral Beach

#3 Aloft #69 Ocean Dream

#16 RIU Cancun #6 RIU Palace

#77 Sunset Royal

#30 Le Blanc

#54 Dreams Sands

#11 Fiesta Americana Villas

What are the drivers of hotel relevance?

Credit: Matteo Colombo/Getty Images

Force plots reveal the effect of every feature

Grand Fiesta: High ranked because it's popular with decent rewards



Sunset Royal: Low ranked because no history, reputation, and a high price



Every data field is credited with its contribution to the final prediction. This credit is the data field's **SHAP value**.

Introduction

Evaluation

Interpretation

Machine Learning

Deployment

Violin plots summarise aggregate effects



Partial dependence plots for feature interactions **¬**R

Model is able to **adjust the score** based off of how closely the **hotel price** (srq_price_zscore) matches the **user's preferred price level**



Learning

Deployment

Introduction

Modeling search result relevance





Model of choice: gradient boosting tree



We can continue gradient boosting ad infinitum



FR

Predict which hotel is "preferred" by the customer

- Unconsidere	o Irrelev	ant <u>1</u> Details	2 Payment	3 Purchased
1. Mgm Grand	2. Aria	3. Ti - Treasure Island	4. The Venetian®	5. Hard Rock
6. Paris	7. The Palazzo	8. The Cosmopolitan	9. Luxor	10. Sls
11. Tropicana	12. Westgate	13. The Linq	14. Bally's	15. Circus Circus
16. Golden Nugget	17. The Signature	18. El Cortez	19. Renaissance	20. Polo Towers
21. Waldorf Astoria	22. Palms Casino Resort	23. Red Rock	24. Nomad Las Vegas	25. Downtown Grand

A result's label is determined based on its **funnel stage**

Introduction

Interpretation

Machine Learning ٦R

Deployment

Predict which hotel is "preferred" by the customer

- Unconsidere	ed o Irrelev	ant <u>1</u> Details	2 Payment	3 Purchased
1. Mgm Grand	2. Aria	3. Ti - Treasure Island	4. The Venetian®	5. Hard Rock
6. Paris	7. The Palazzo	8. The Cosmopolitan	9. Luxor	10. Sls
11. Tropicana	12. Westgate	13. The Linq	14. Bally's	15. Circus Circus
16. Golden Nugget	17. The Signature	18. El Cortez	19. Renaissance	20. Polo Towers
21. Waldorf Astoria	22. Palms Casino Resort	23. Red Rock	24. Nomad Las Vegas	25. Downtown Grand

Neighbors of positively labeled results were likely **considered**

Introduction

Interpretation





ΞR

Predict which hotel is "preferred" by the customer

- Unconsidere	ed o Irrelev	ant <u>1</u> Details	2 Payment	3 Purchased
1. Mgm Grand	2. Aria	3. Ti - Treasure Island	4. The Venetian®	5. Hard Rock
6. Paris	7. The Palazzo	8. The Cosmopolitan	9. Luxor	10. Sls
11. Tropicana	12. Westgate	13. The Linq	14. Bally's	15. Circus Circus
16. Golden Nugget	17. The Signature	18. El Cortez	19. Renaissance	20. Polo Towers
21. Waldorf Astoria	22. Palms Casino Resort	23. Red Rock	24. Nomad Las Vegas	25. Downtown Grand

Results **unconsidered** by the user are left out of the dataset

Introduction

Interpretation





٦R

$$\lambda_i = \sum_j [P_{\text{true}}(x_i \gg x_j) - P_{\text{predicted}}(x_i \gg x_j)] |\Delta_{NDCG}|$$

- The model is graded based on how far the predicted preference probability differs from the true preference probability

Evaluation

Introduction

- **NDCG** is a metric which penalizes errors made high in the ranking more than mistakes made further down
- When we gradient boost with this error function, we are using a **Learning to Rank** model called **LambdaMART**

Interpretation

Machine

Learning

Deployment

Map score differences to preference probabilities

$$P_{\text{predicted}}(x_i \gg x_j) = \text{sigmoid}(x_i - x_j) = \frac{e^{x_i}}{e^{x_i} + e^{x_j}}$$



٦R

~ ~

Converting scores to probabilities for profit reranking **¬**R

$$\frac{\operatorname{sigmoid}(x) \to \operatorname{softmax}(\{x_1, \dots, x_i, \dots, x_n\})}{e^x + e^0} \xrightarrow{e^{x_i}} e^{x_1} + \dots + e^{x_i} + \dots + e^{x_n}}$$

This **softmax probability** represents the chance that result X_i from the set {X_j} is picked, given the user picks one result

With these probabilities, we can rerank results and maximize expected profit!

(Actually, since the softmax normalization is constant across all results in the request, simply multiplying profit by the exponentiated scores achieves the profit-maximizing sort,)

Interpretation

Machine

_earning

Deployment

Introduction

Evaluation

Deployment

0000 0000 00000000000 0.1

C.

Credit: Sergly Serdyuk/IStockPhoto/Thinkstock





















Quality assurance at every step



Data integrity

- Nulls
- Duplications
- Ranges
- Similarity to prod logs

Model quality

- NDCG scores
- SHAP signs
 - Correlations with previous models

Infrastructure

- Unit tests
- Batch load simulations

Machine

Learning

Performance

- MBPR
- Prediction histograms
- Data integrity

Deployment

Introduction

Engineering

Reducing Latency Encapsulability QA and monitoring Broader model ecosystem



Theory

Embeddings as features Alternate loss function Debiasing by rank Supporting dynamic reranking

Thanks to everyone who contributed!









Eric He Research and Analysis

Renyu Zhang Architecture and QA

Gary Ng Planning and Review



The full list of mathematical equations

- 1. Map score differences to preference probabilities $P(x_i \triangleleft x_j) = \sigma(I-J) = \frac{e^I}{e^I + e^J}$
- 2. Loss function adds NDCG-reweighting and rank debiasing

$$\lambda_{x_i} = \sum_{x_j} [P_{\text{true}}(x_i \triangleleft x_j) - P_{\text{predicted}}(x_i \triangleleft x_j)] * \frac{|\Delta \text{NDCG}|}{p^+(\text{rank}(x_i)) * p^-(\text{rank}(x_j))}$$

$$DCG = \sum_{i} \frac{2^{rel_{i}} - 1}{\log_{2}(i+1)} \qquad p_{raw}^{+}(1) = \sum_{q} \sum_{j} \frac{L(1,j)}{p^{-*}(j)} \qquad p^{+}(i) = \frac{1}{p_{raw}^{+}(1)} \sum_{q} \sum_{j} \frac{L(i,j)}{p^{-*}(j)}$$
$$NDCG = \frac{DCG}{IDCG} \qquad p_{raw}^{-}(1) = \sum_{q} \sum_{j} \frac{L(1,j)}{p^{+*}(j)} \qquad p^{-}(i) = \frac{1}{p_{raw}^{-}(1)} \sum_{q} \sum_{j} \frac{L(i,j)}{p^{+*}(j)}$$

3. Multiply exponentiated relevance score by profitability metric

Ranking Score =
$$Relevance * \frac{Profit}{\sqrt{Revenue}}$$

Updated architecture

