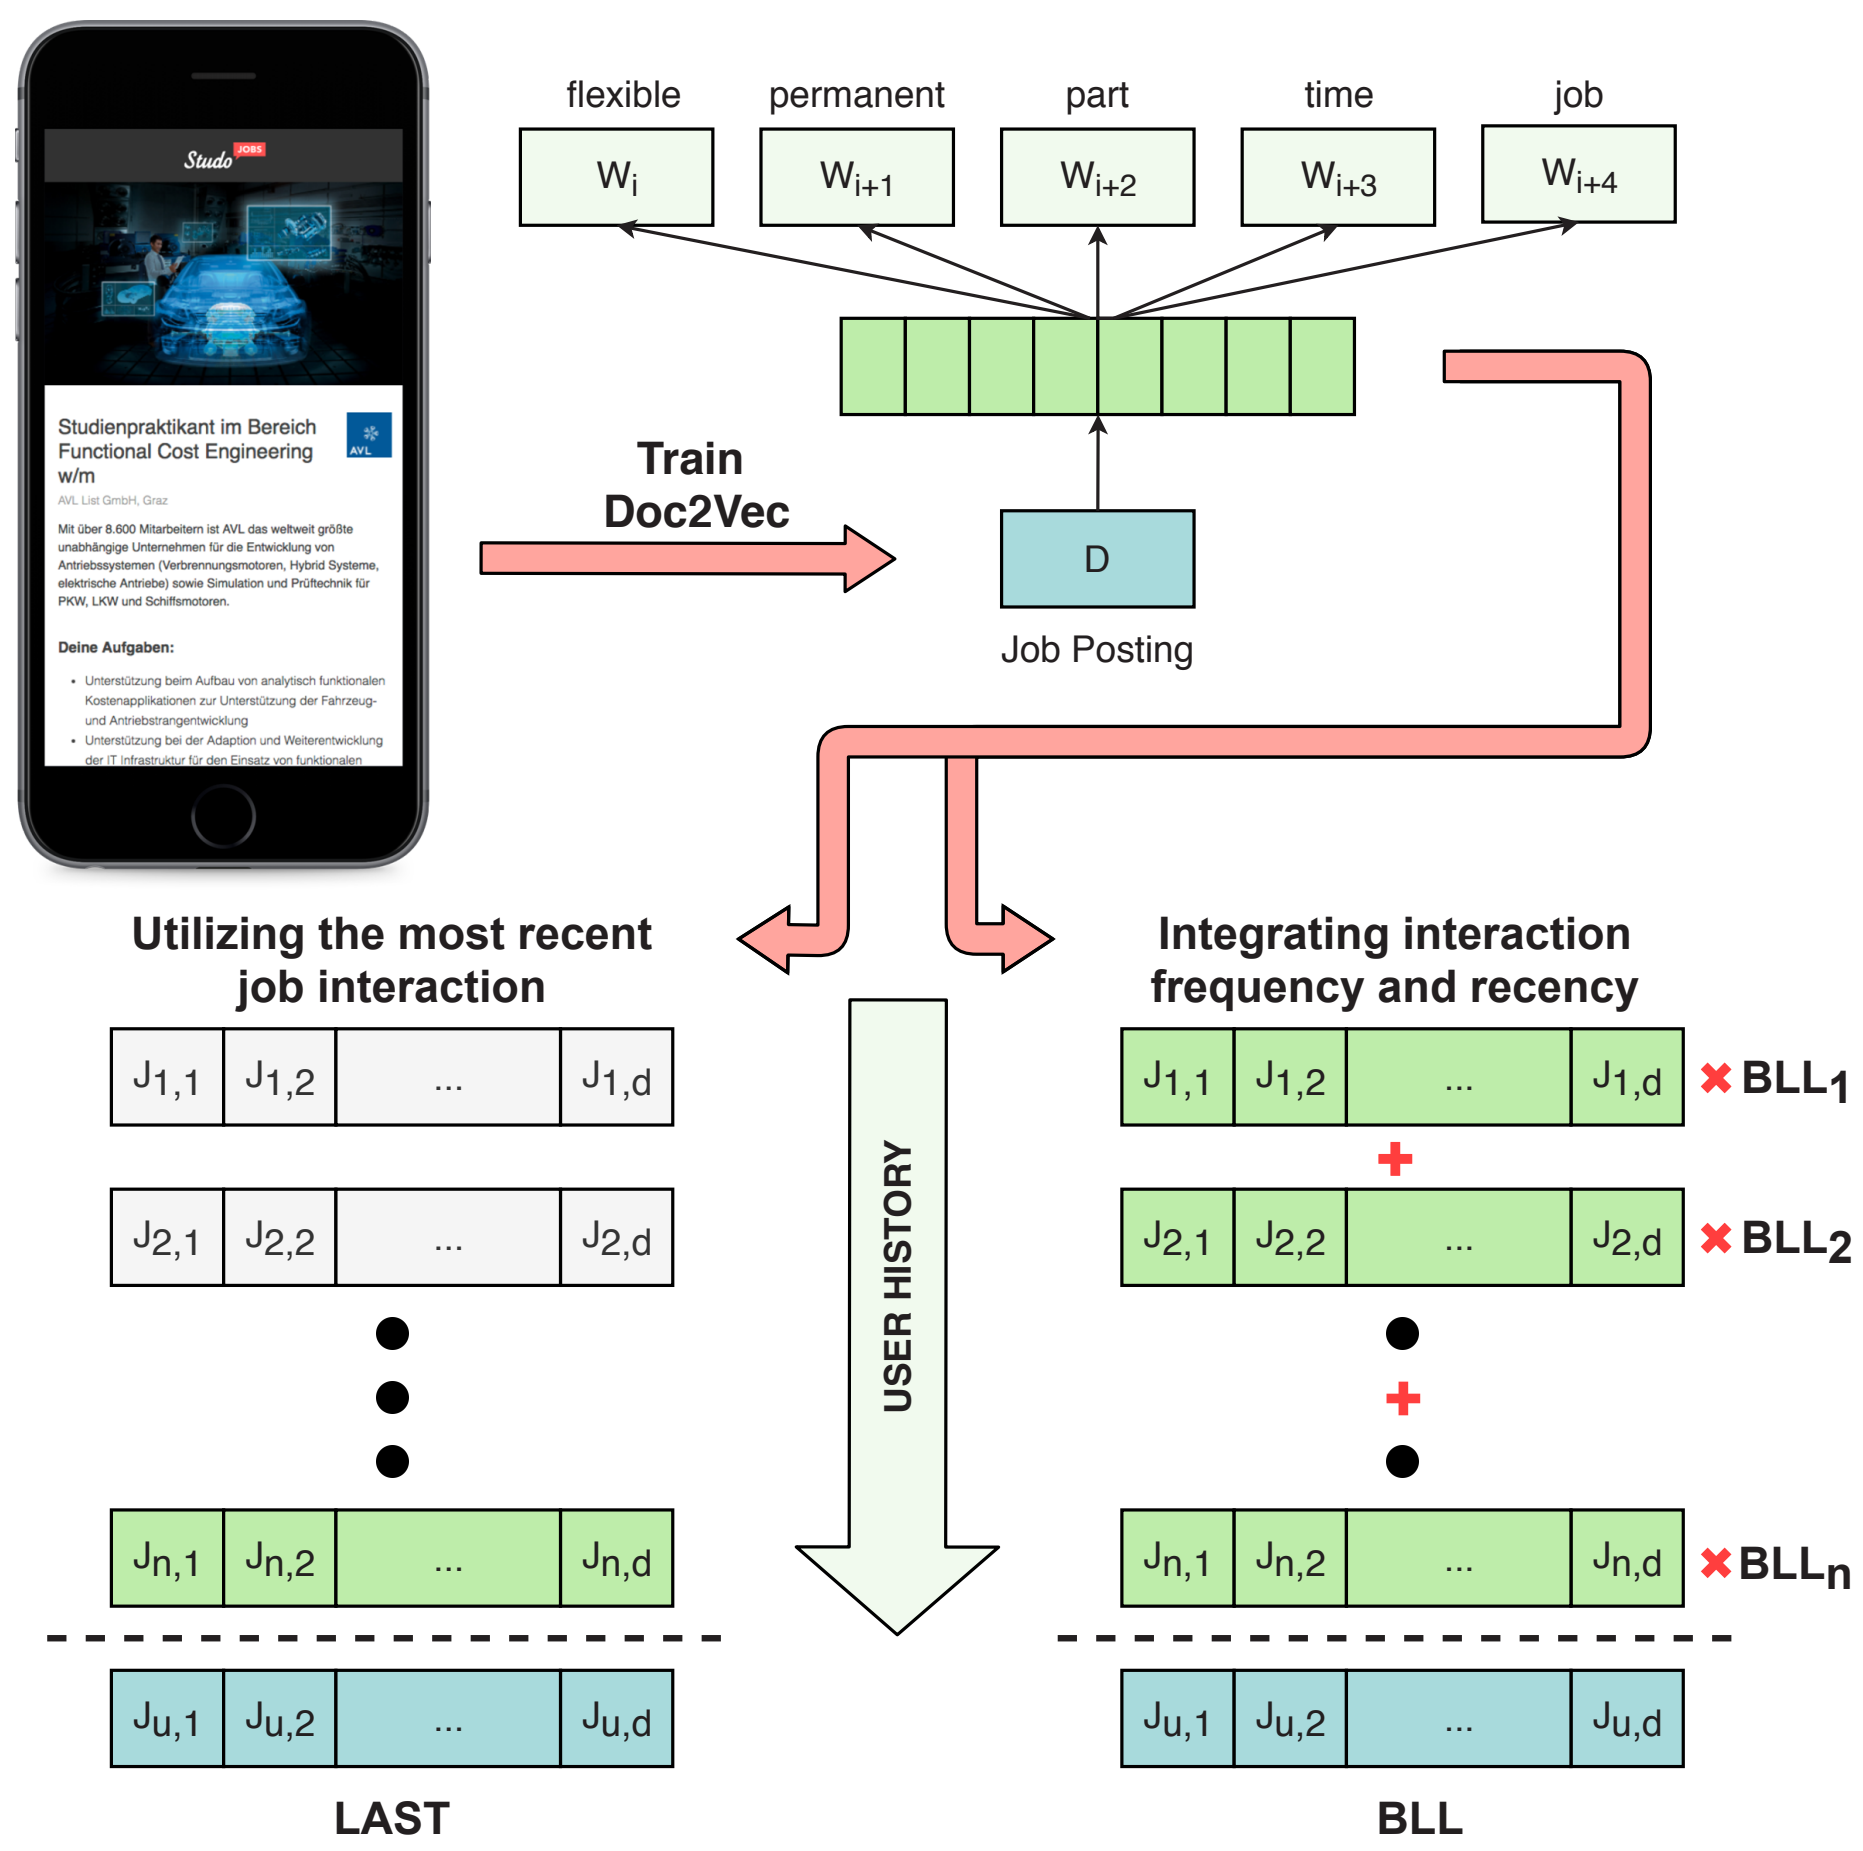


# SHOULD WE EMBED?

## A STUDY ON THE ONLINE PERFORMANCE OF UTILIZING EMBEDDINGS FOR REAL-TIME JOB RECOMMENDATIONS

### STUDY SETUP

**Impact of embeddings** analysed under real-time constraints on the Austrian job platform Studo Jobs<sup>a</sup>.



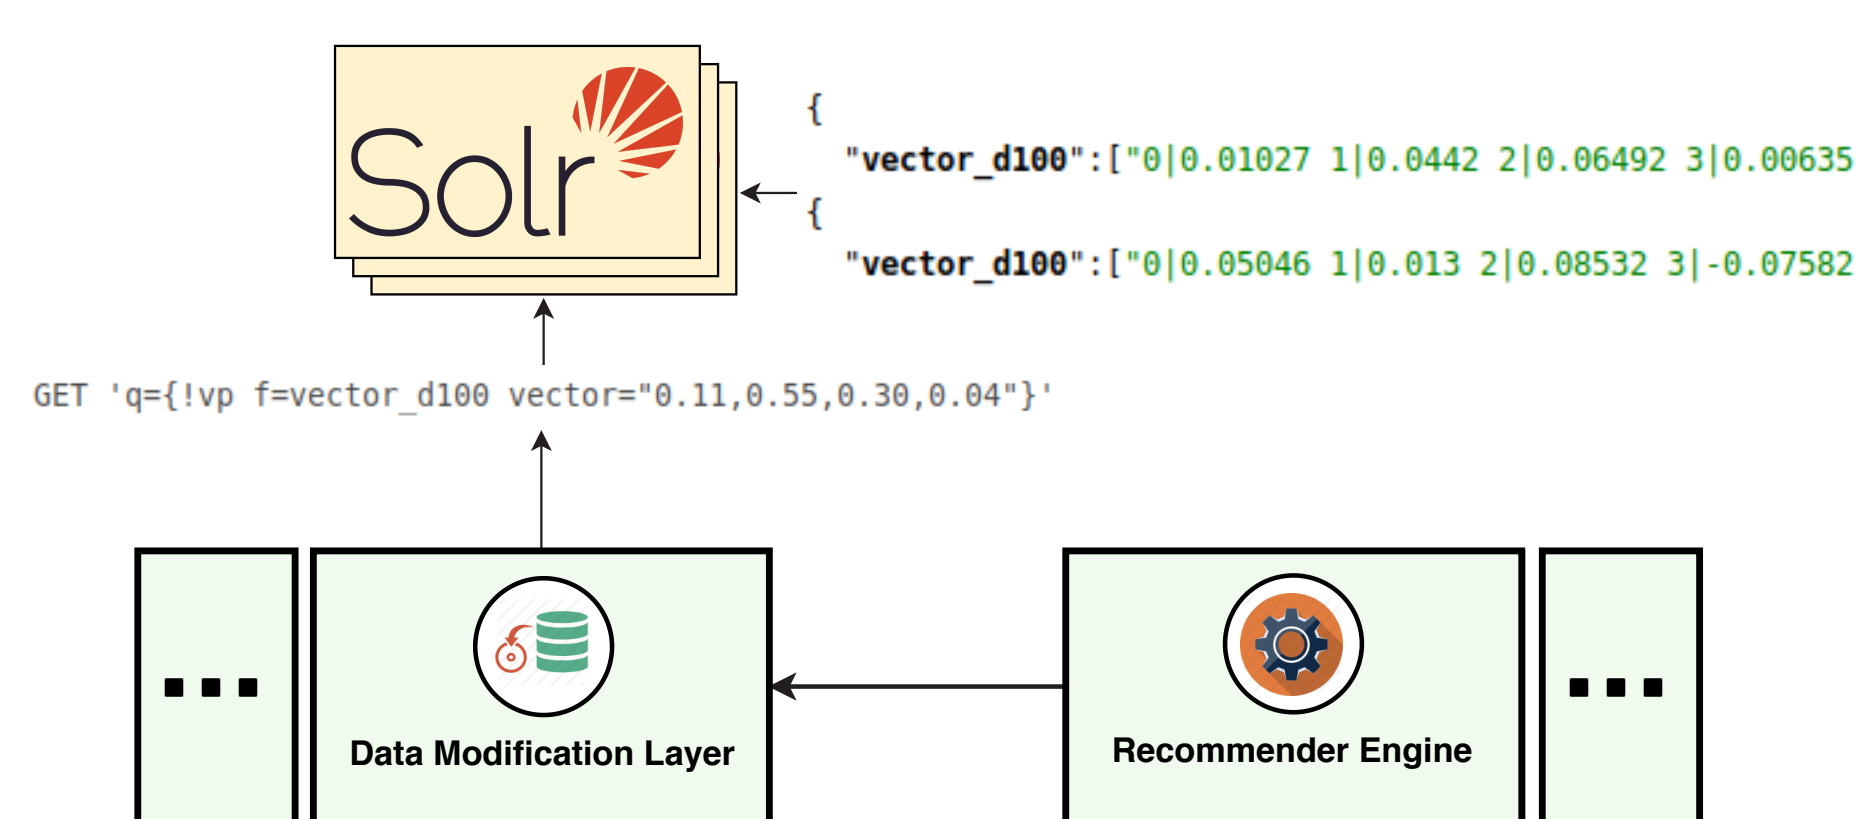
**LAST.** A natural way to utilize embeddings is to consider the most recent job posting and recommend top-*k* similar jobs.

**BLL.** Accounts simultaneously for both **frequency** and **recency**. Browsing behaviour for a reference vector is modelled by:

$$BLL_{u,j} = \ln\left(\sum_{i=1}^n (TS_{ref} - TS_{j,i})^{-d}\right)$$

where *d* represents the time-dependent decay parameter.

**Adapting for real-time job recommendations.** Response times need to be below 100-200 milliseconds. Embeddings can be stored as payloads in Lucene and Cosine similarity can be used for retrieval.



<sup>a</sup>Superseded by the Talto career platform: <https://talto.com>

### CONCLUSION

Using embeddings **significantly improves** the CTR and runtime performance for recommending **similar jobs**.

Factors of **recency** seem to be especially **influential** in this setting.

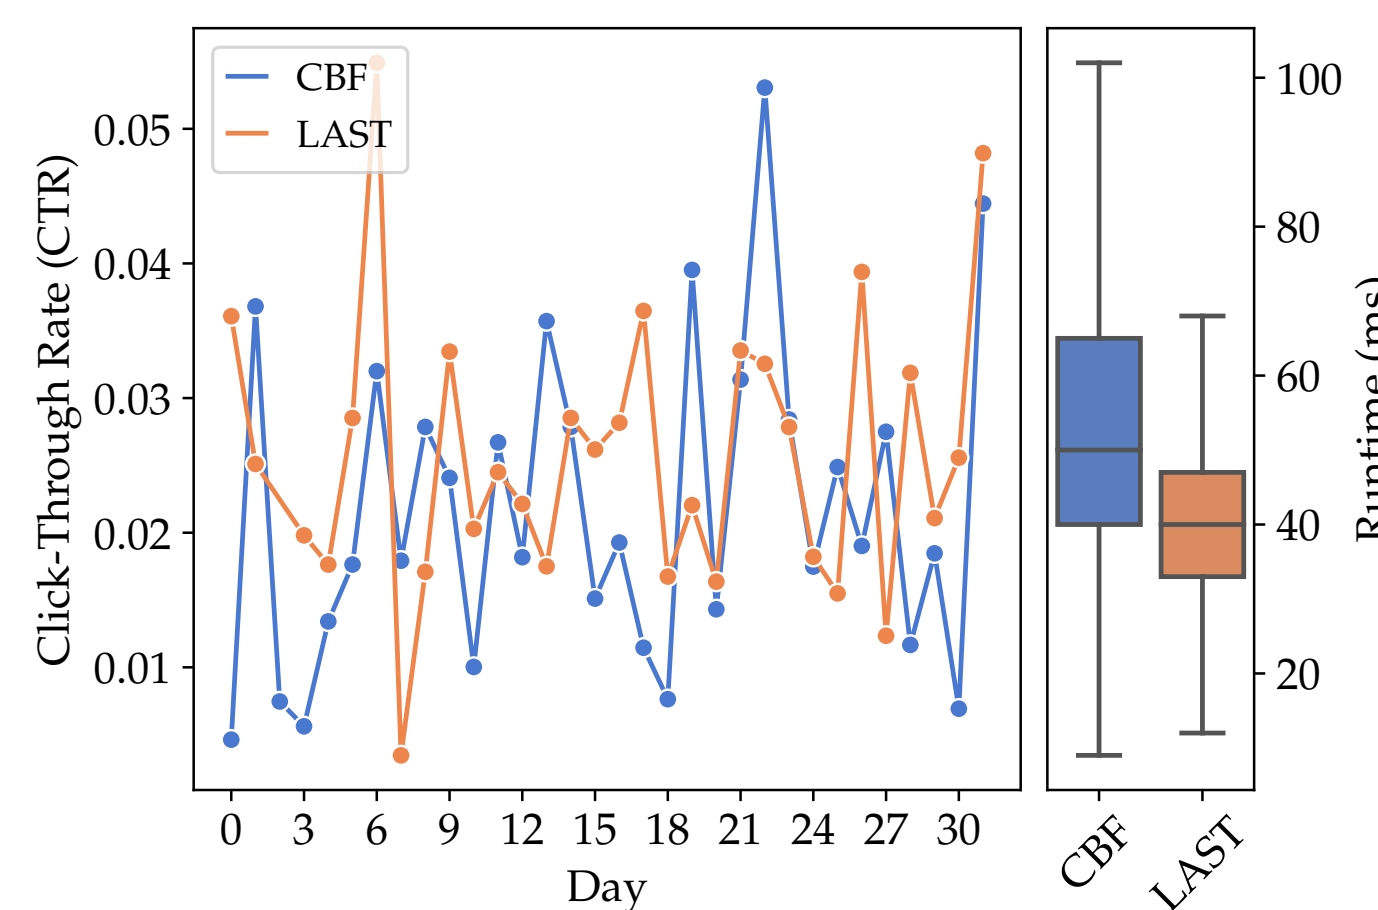
By combining embeddings based on the factors of **frequency** and **recency**, we can further **enhance** the online performance when personalizing the **homepage**.

### SIMILAR JOBS

**Preliminary Analysis: Embedding size.** In the case of Studo, embeddings larger than 100 **did not contribute to a higher CTR**, but **did increase** the overall runtime performance.

**Baseline: Content-Based Filtering (CBF).** Popular in many systems for recommending similar items. Easily adaptable when recommendations need to be served in real-time.

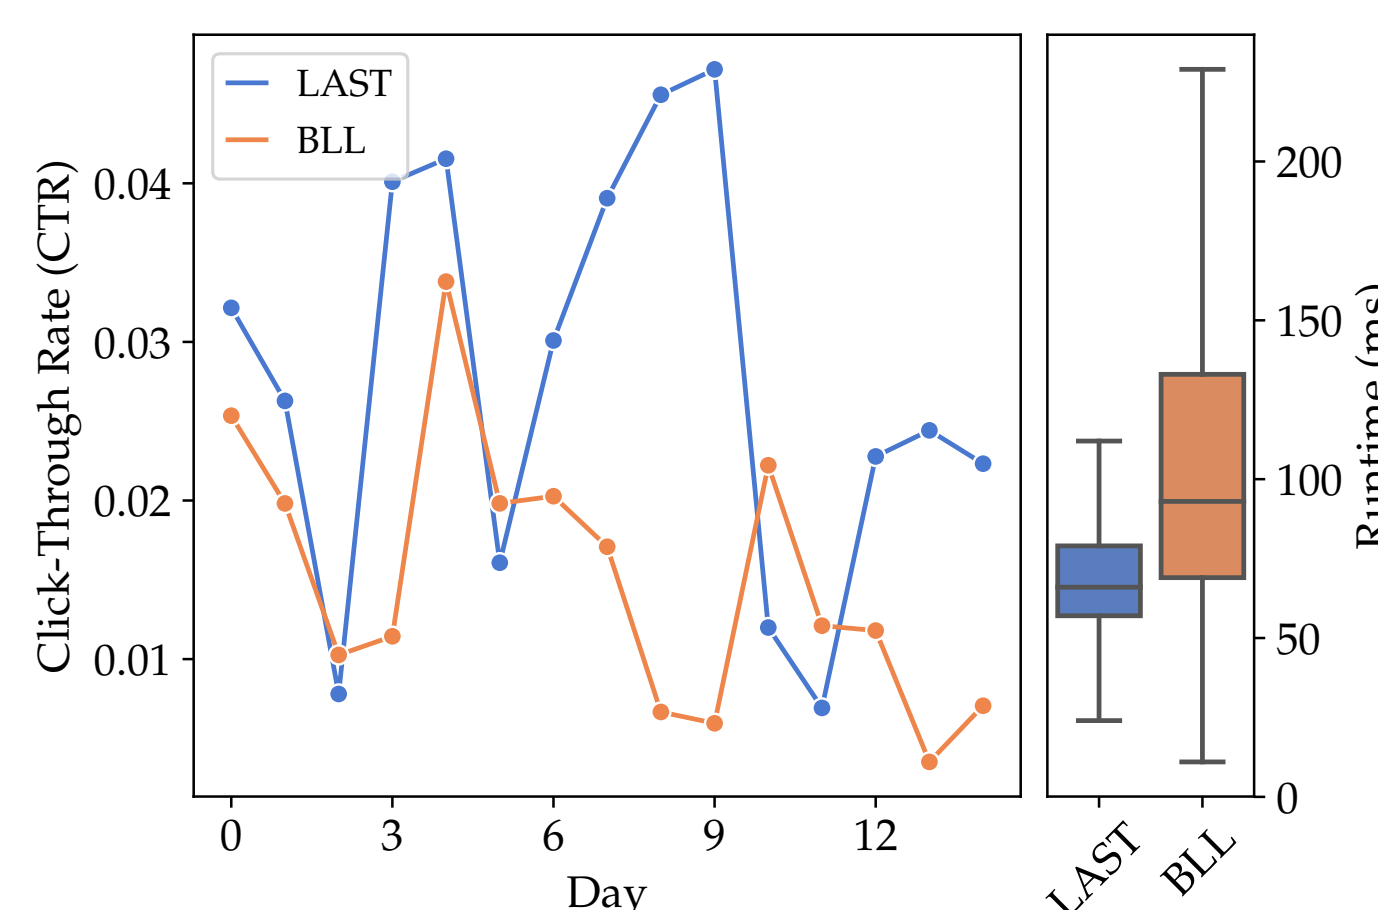
#### Impact of embeddings



# Days: 32  
# Distinct Users: 8, 576  
# Recommendation Requests: 31, 968

Approach	CTR	↑	Runtime (ms)	↓
CBF	0.0194	18.04%	51	23.53%
LAST	<b>0.0229*</b>		<b>39**</b>	

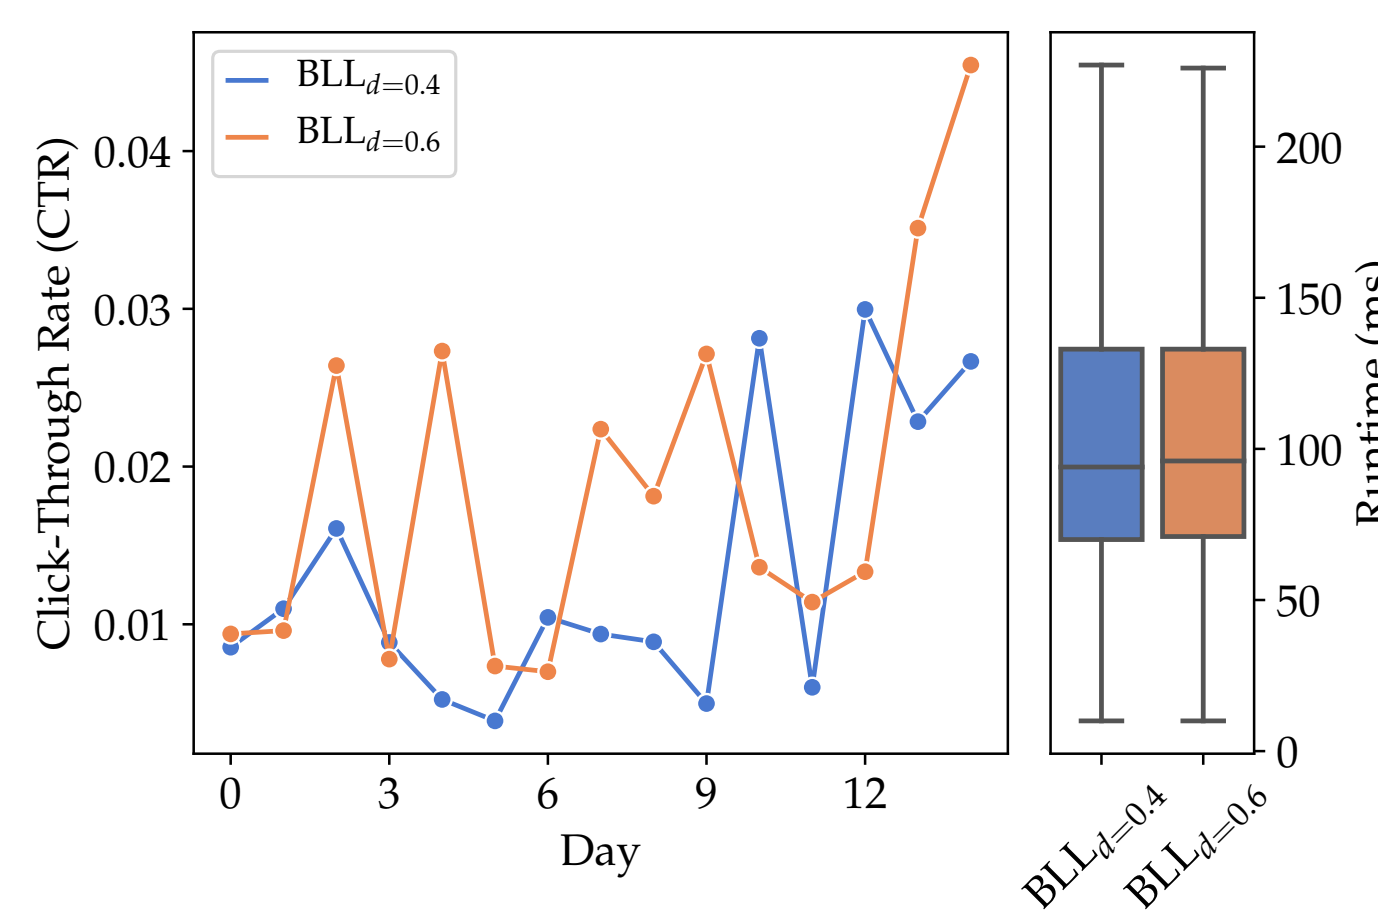
#### Influence of frequency and recency



# Days: 15  
# Distinct Users: 4, 715  
# Recommendation Requests: 18, 464

Approach	CTR	↑	Runtime (ms)	↓
LAST	<b>0.0249**</b>	75.35%	<b>67**</b>	28.72%
BLL	0.0142		94	

#### Merit of recency




# Days: 15  
# Distinct Users: 3, 375  
# Recommendation Requests: 11, 992

Approach	CTR	↑	Runtime (ms)	↓
BLL <sub>d=0.6</sub>	<b>0.0174*</b>	35.94%	97	2.06%
BLL <sub>d=0.4</sub>	0.0128		<b>95</b>	

### HOMEPAGE

**Baseline: Collaborative Filtering (CF).** One of the most explored and utilized techniques for personalizing a system in real-time. To account for cold-start users, we recommend the most popular job postings as a fallback. To provide recommendations in real-time, the inverted-index structure available in the Apache Solr search engine is used to find the *k*-nearest neighbors using the Cosine similarity metric.

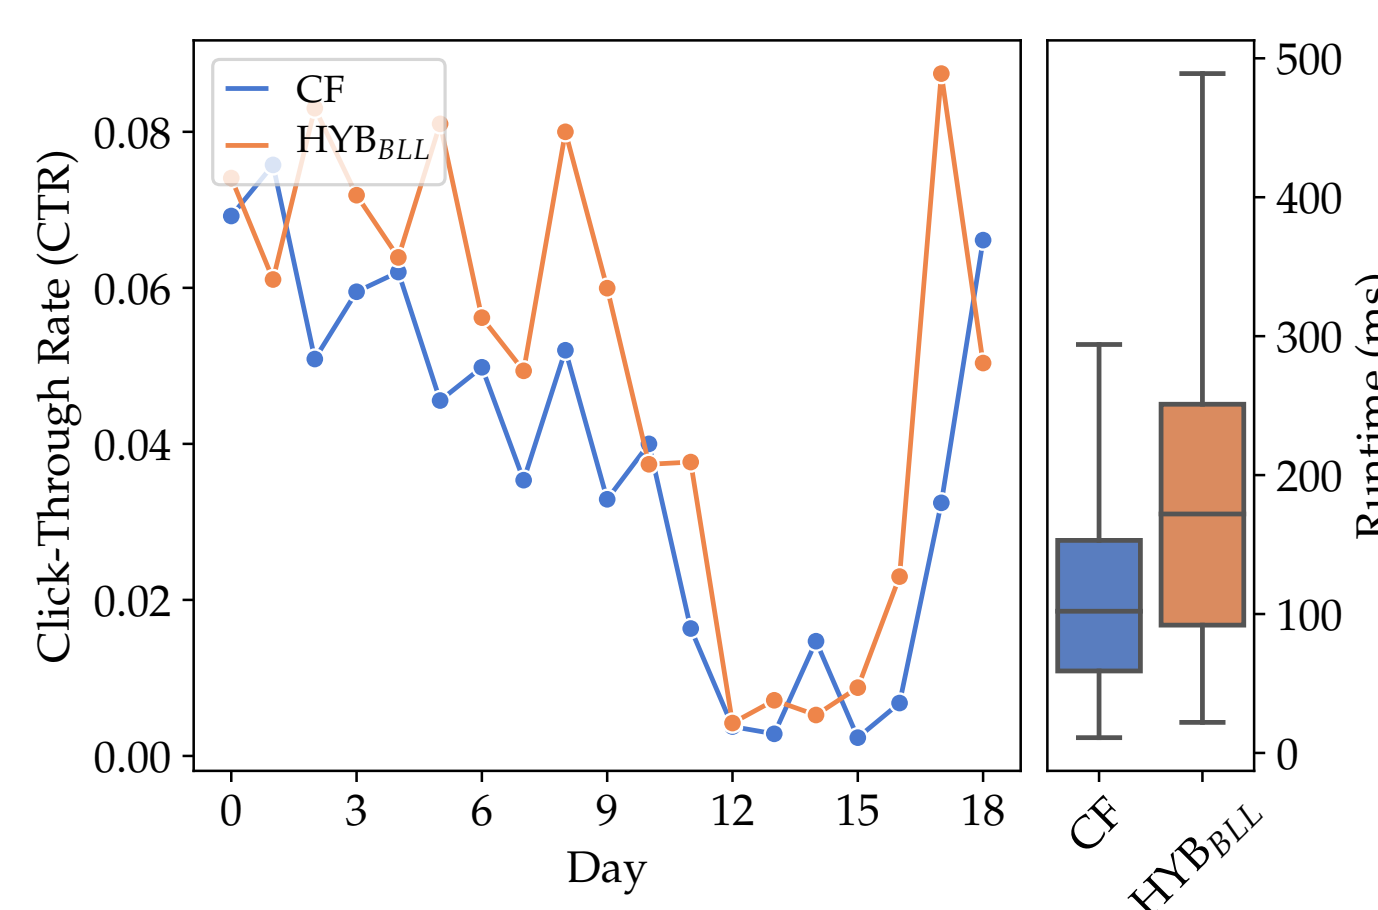
#### Influence of frequency and recency



# Days: 25  
# Distinct Users: 9, 620  
# Recommendation Requests: 26, 334

Approach	CTR	↑	Runtime (ms)	↓
BLL	<b>0.0671*</b>	15.69%	<b>114**</b>	13.64%
CF	0.0580		132	

#### Combining frequency and recency



# Days: 19  
# Distinct Users: 9, 313  
# Recommendation Requests: 24, 907

Approach	CTR	↑	Runtime (ms)	↓
HYB <sub>BLL</sub>	<b>0.0471**</b>	33.05%	172	38.37%
CF	0.0354		<b>106**</b>	