wayfair

Association for Computing Machinery

Contextual Product Recommendation at Wayfair

Jeffrey Mei, Ph. D.
Senior Data Scientist

CARS Workshop 2022

Wayfair - All Things Home

0.

Context

Static vs. Dynamic Context at Wayfair

Static

"A set of observable attributes that are known a priori" that influence customer behavior

Customer location

Customer gender

Day of week

Whether a sale is on

Dynamic

"A set of conditions under which an activity occurs … where the activity gives rise to context and the context influences activity"

"Taste" (may change between, or even within, browse sessions)

Intent (does the user intend to order, or just browse?)

1.

Examples of (Static) Context

Use location to predict what category a customer will browse

California buys different stuff than North Dakota

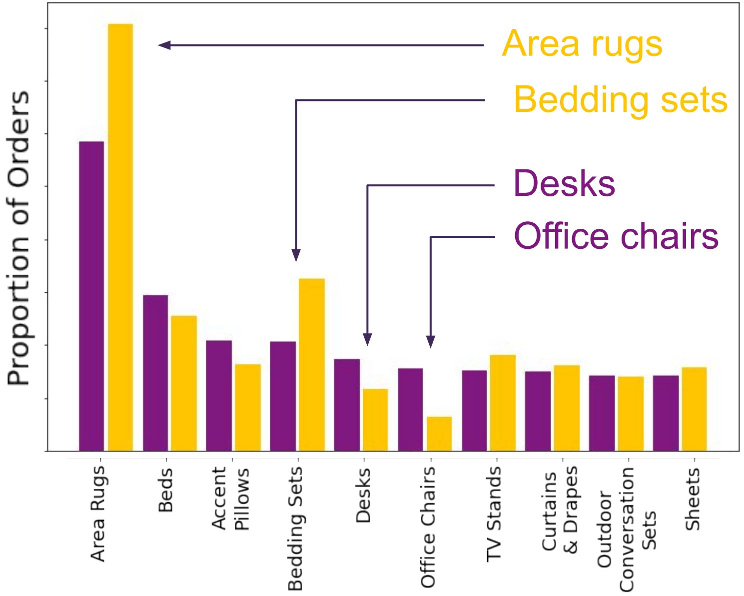

Customer Location

- Use their **location** to factor shipping cost into top recs

- Win for both **suppliers** (shorter shipping distance = lower likelihood of damage) and for **customers** (receive order faster)

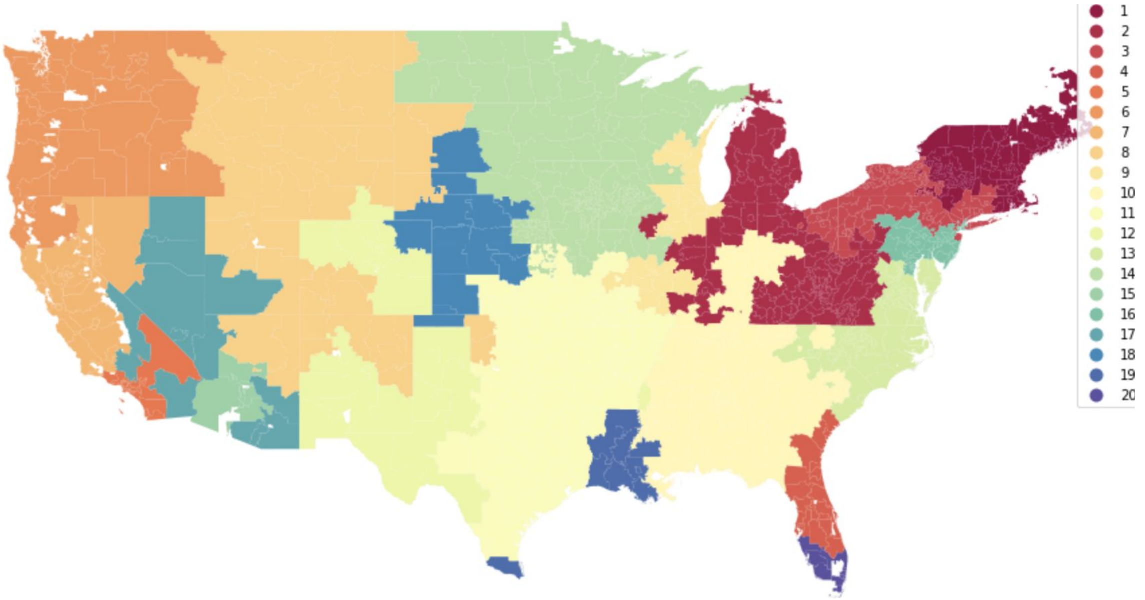

Geo-specific models

- Wayfair serves the US, Canada, UK and DE markets

- Each of these have slightly different catalogs, and also different customer behavior, so we train a specific model for each country

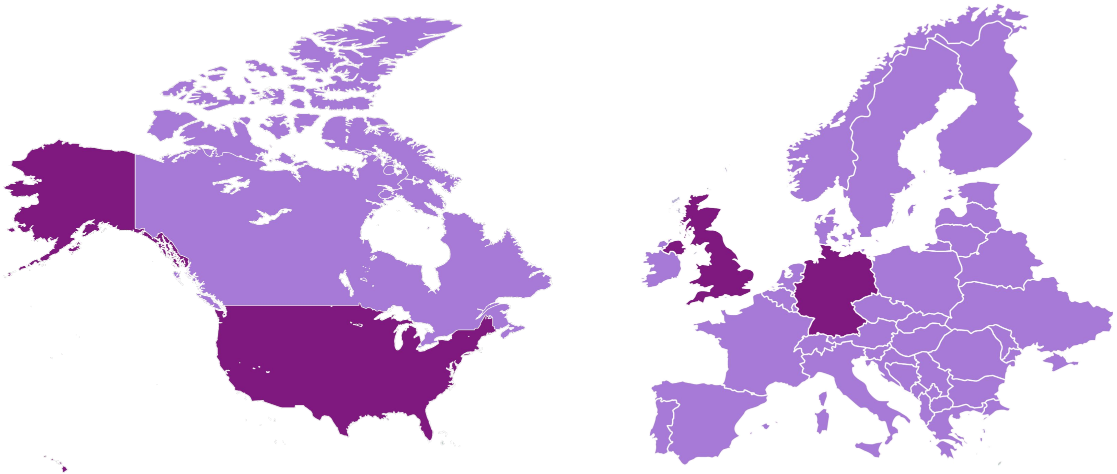

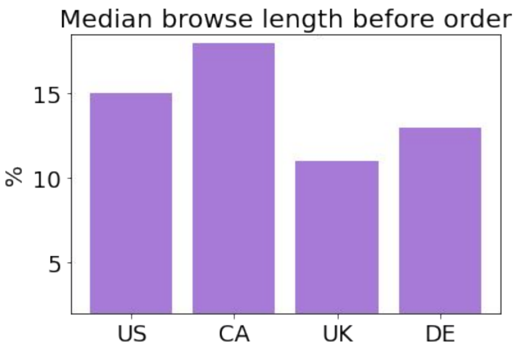

Customer Consideration Cycle

- Customers in all locations generally spend **twice** as much time browsing **indoor furniture** before ordering than for **outdoor furniture**

- Our North American customers typically browse **50% longer** than our European customers before finalizing their order

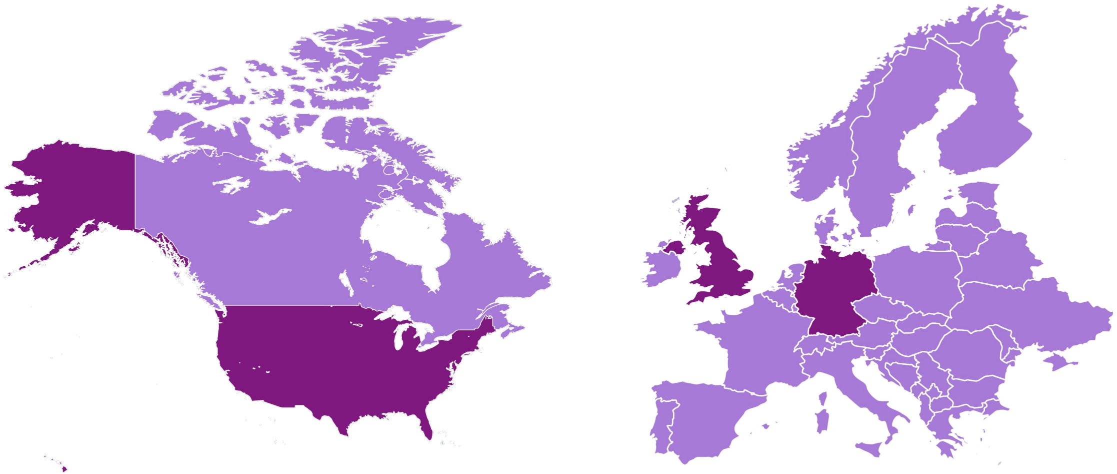

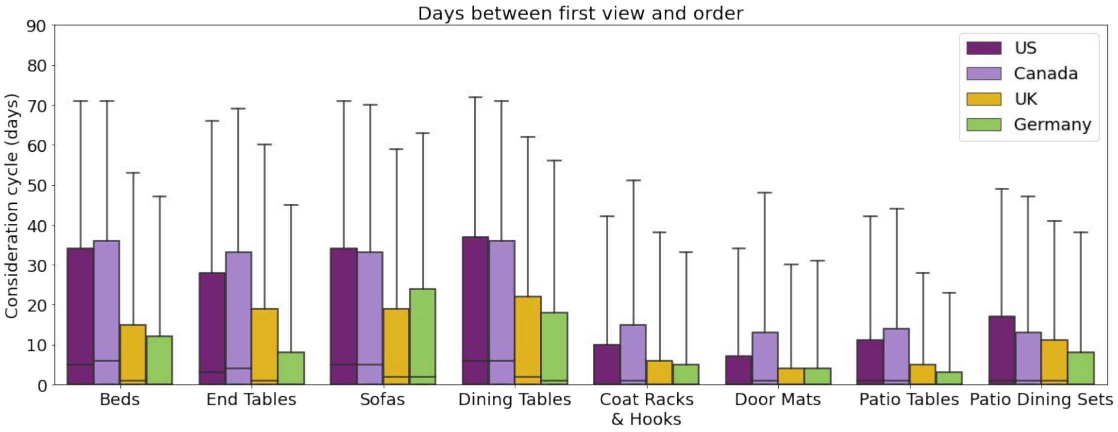

Customer Consideration Cycle

- Customers will spend more than **twice** as much time choosing their **own** bed vs. their **teenager**'s bed

- In the UK, customers spend **equal** time on both; in Germany, customers spend **virtually no time** choosing their teenager's bed

**Disclaimer: draw conclusions about parenting styles at your own risk*

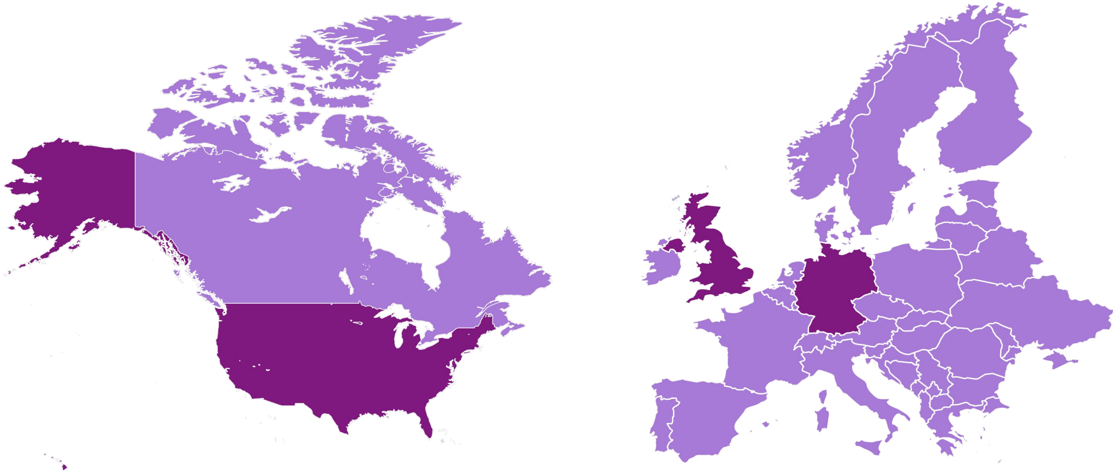

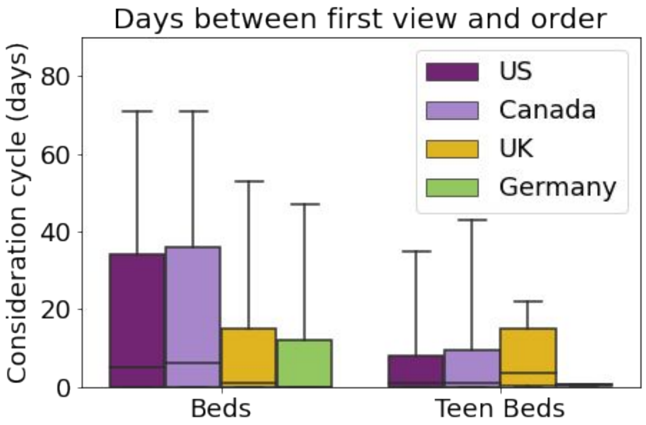

2.

Dynamic Context

What’s in the browse context?

Multi-headed Attention Recommender System

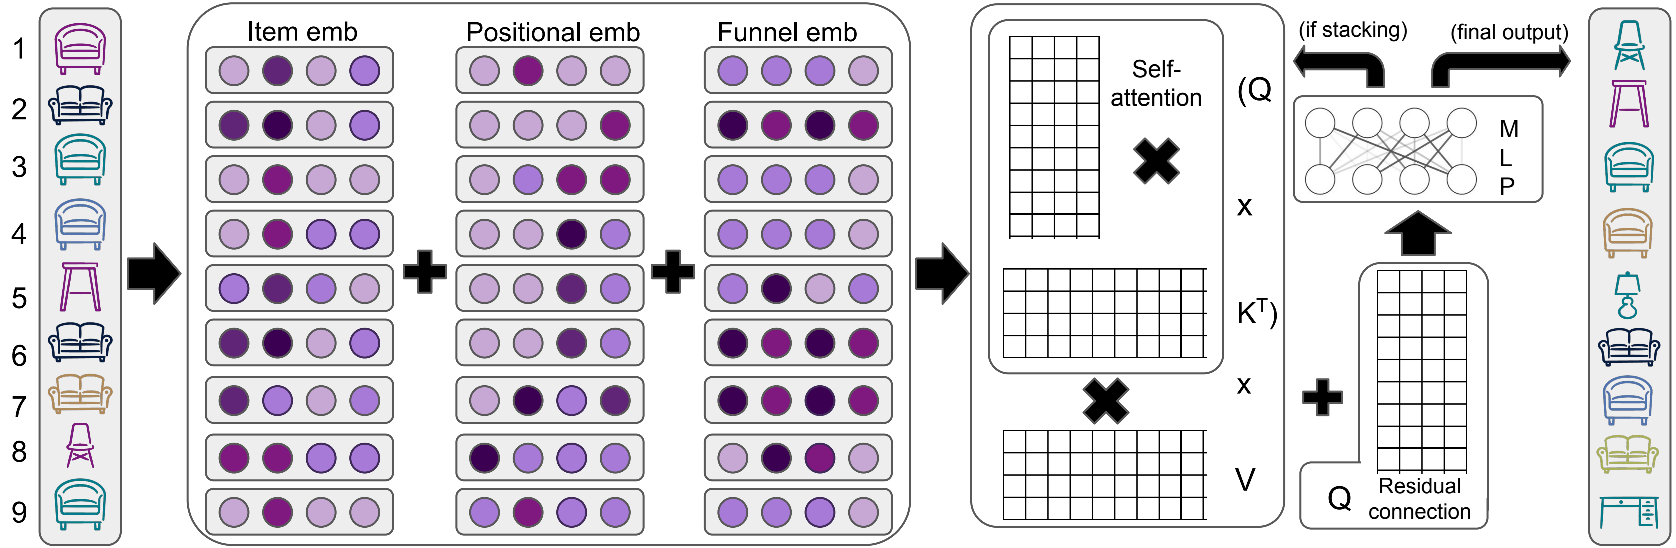

MARS is a transformer network designed to predict the next item that the customer will interact with.

Based off SASRec ([Kang & McAuley 2018])

MARS powers product recommendations

It is trained on **customer-item interactions** - i.e. viewed items, added-to-cart and ordered items (if they exist for a customer), **in the order** they were interacted with.

(link for video)

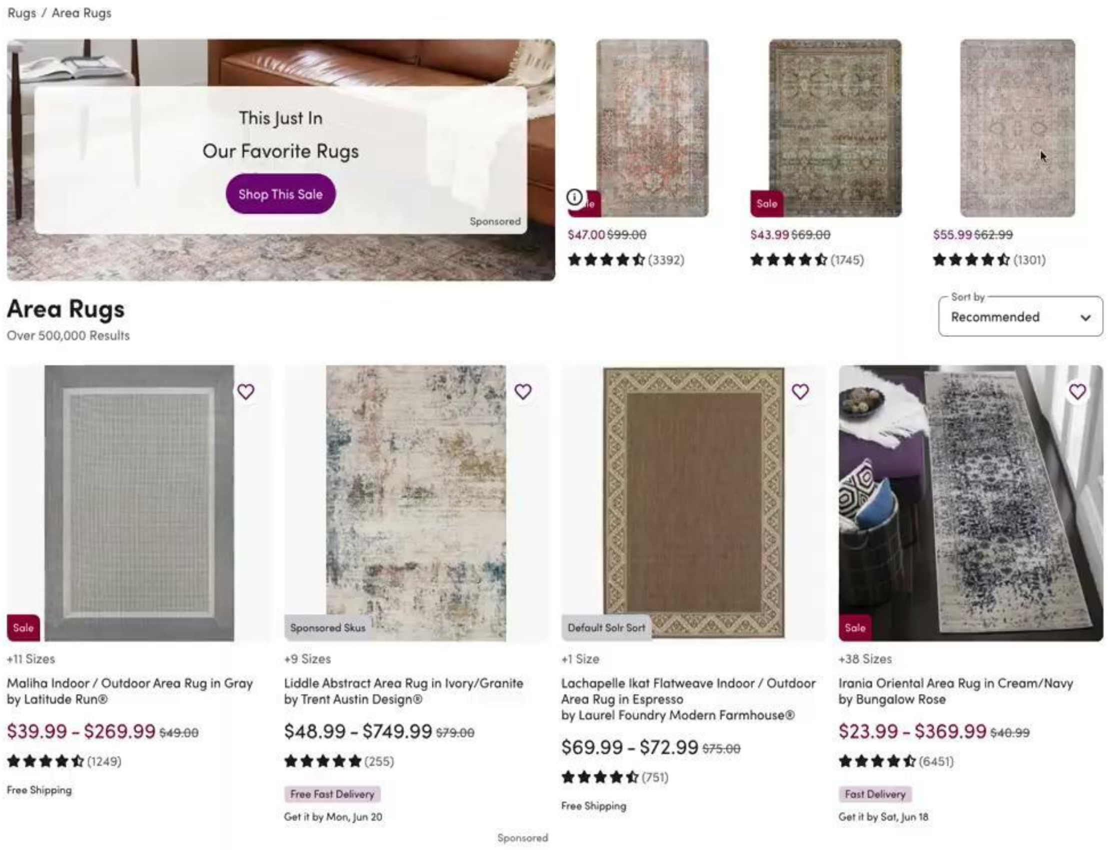

Input - black and white geometric, colorful floral, nautical. Darker = more attention

Using matrix factorization, the majority browse (here, black and white geometric rugs) dominates, and the recommendations are all representative of the **majority** browse **(which is no longer relevant to the customer)**

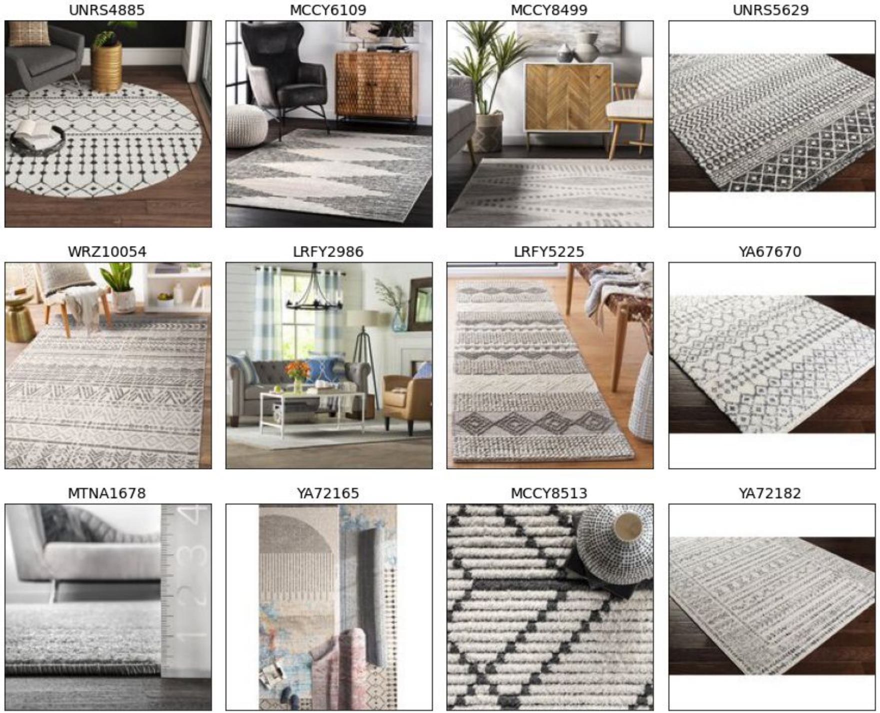

Input - black and white geometric, colorful floral, nautical. Darker = more attention

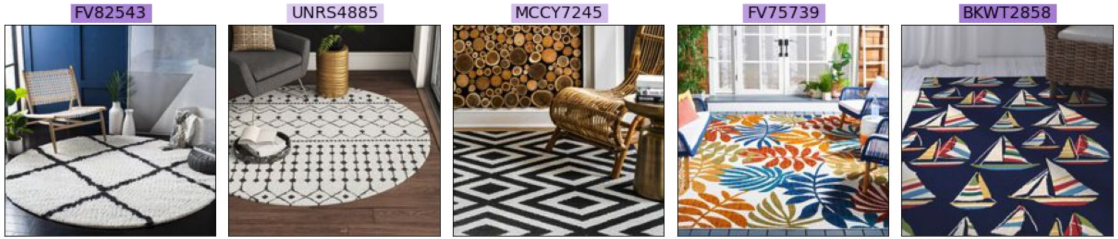

Because MARS is sequential, it is able to adapt to changes in a customer's browse. It pays slightly more attention to the final item (i.e. the **latest customer preference**)

Here, we can see that adding floral and nautical rugs means that the customer is shown "hybrid-style" geometric-nautical, floral-nautical rugs, in addition to more standard nautical and geometric rugs.

This hybridization of style is interesting and we will return to this later…

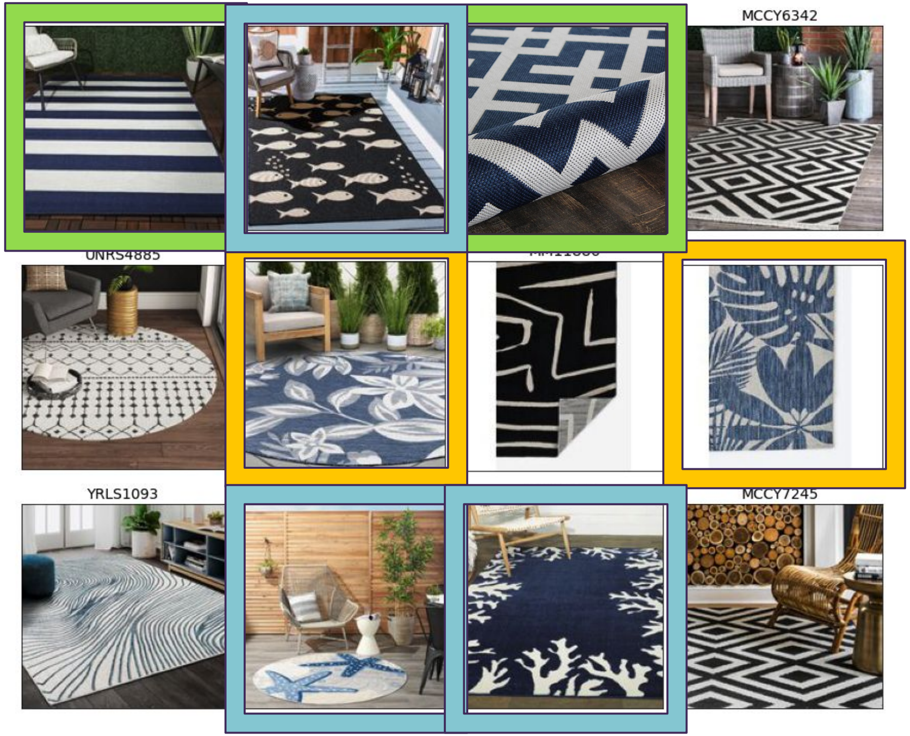

2a.

Position

Paying attention … how?

Learned (MARS)

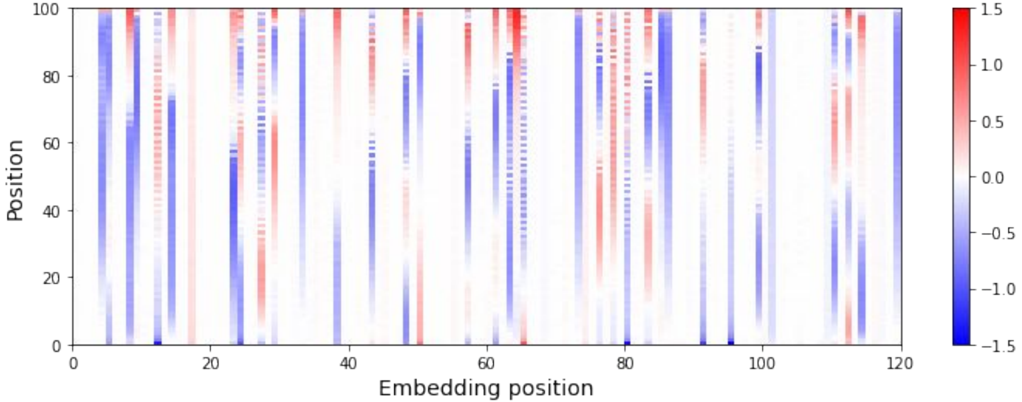

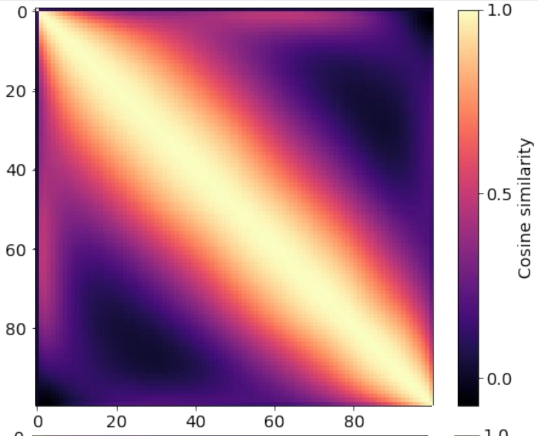

Precomputed (Vaswani 2017)

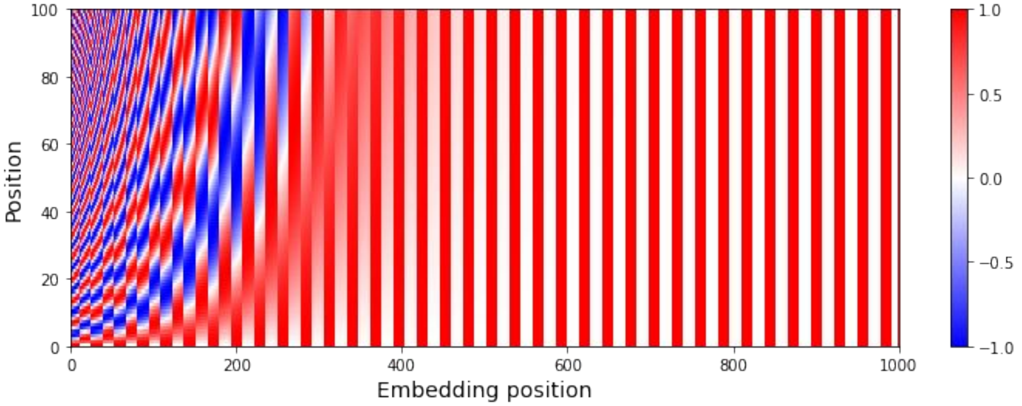

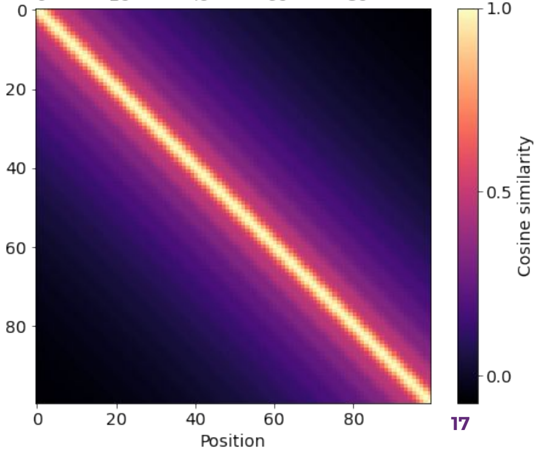

Paying attention … to what?

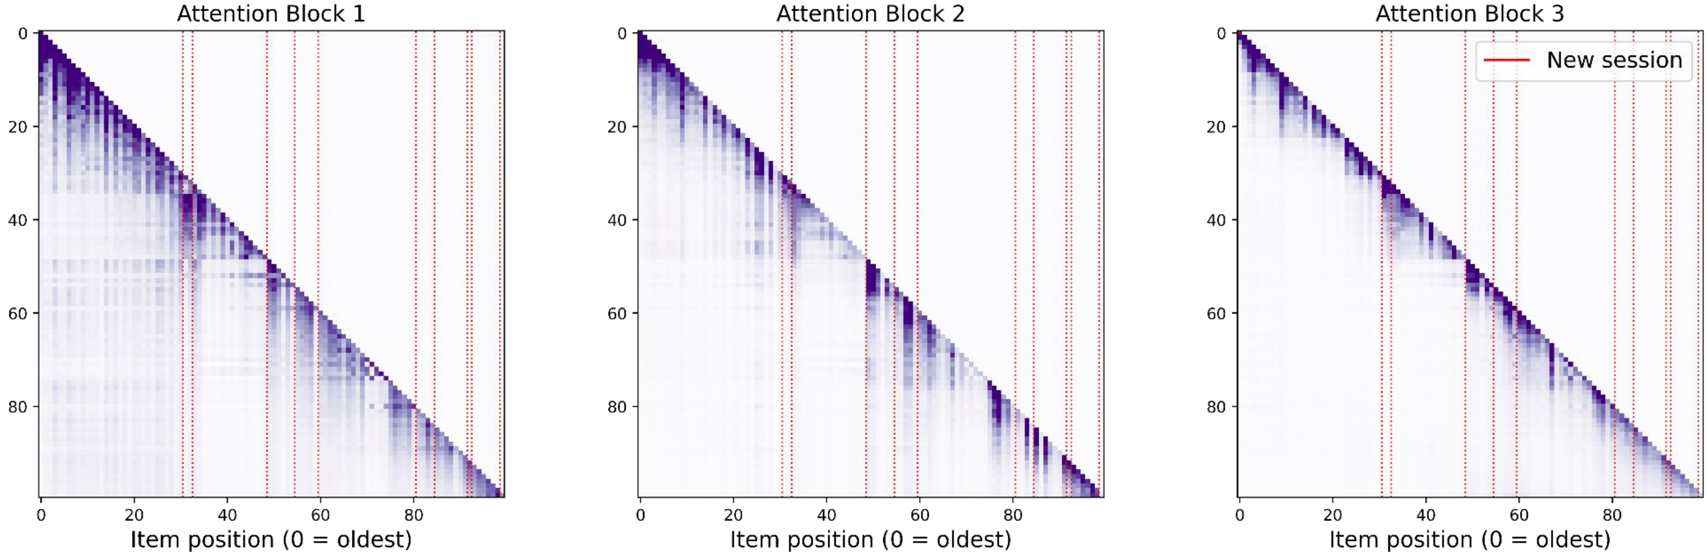

- Earlier blocks attend to longer-range dependencies
- MARS is somewhat session-aware (here defined as a gap of >24 hours) despite having no timestamp information at training
  - Could it be because customers are generally browsing one category per session?

Real-life example

- Many customers browse distinct categories each session. For example, here, the transition from **desks** to **pans** is very clear in the attention weights

- Customers also really like to **view previously-viewed items again** (“**resurfacing**”) - even within the same session

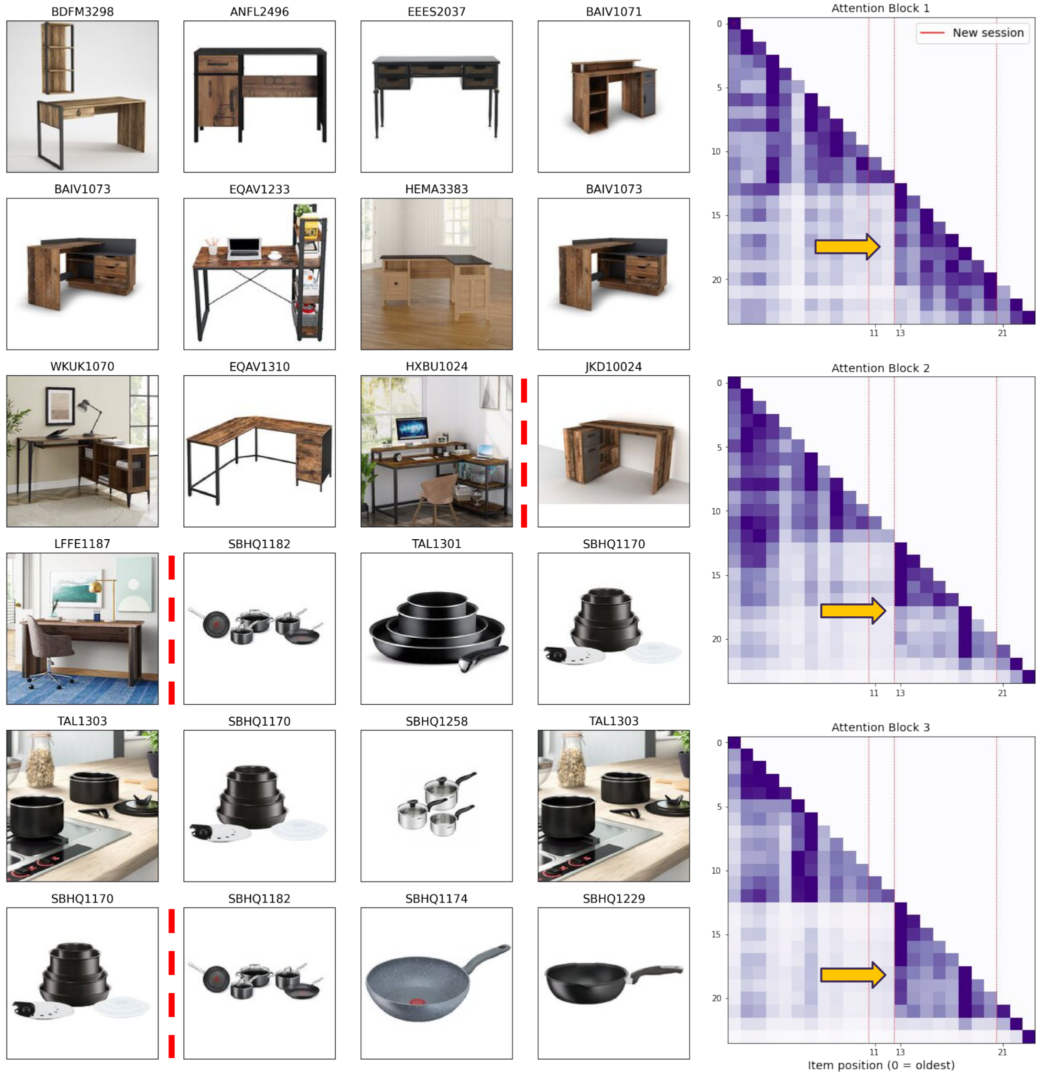

Artificial Example

- No link between sleeper sofas and jute sofas, weak link between bestselling and convertible sofas

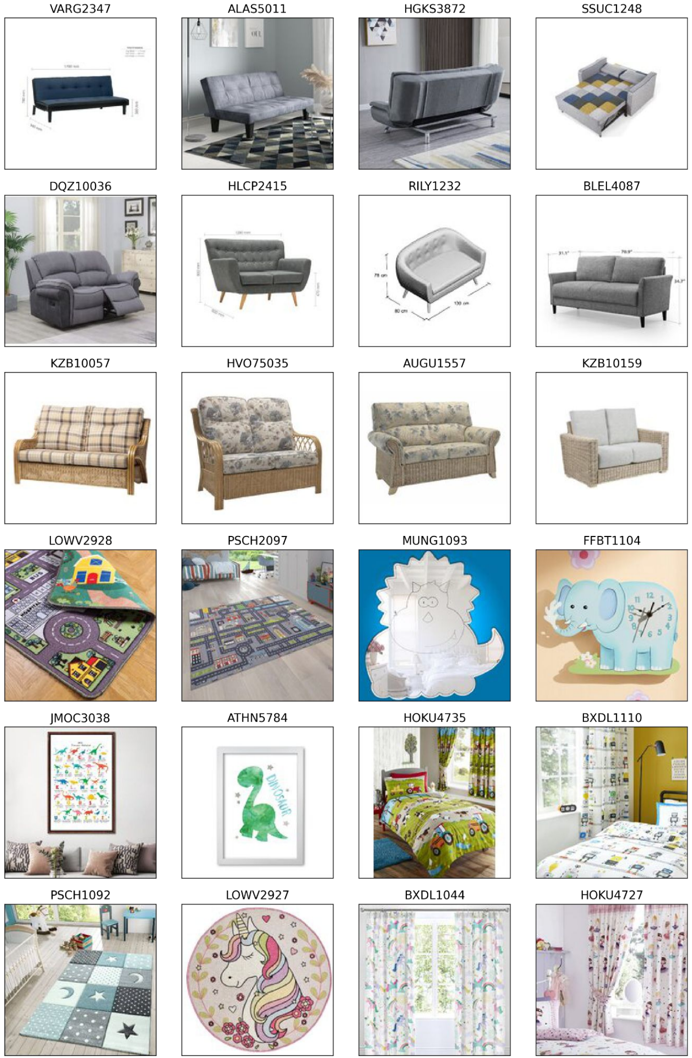

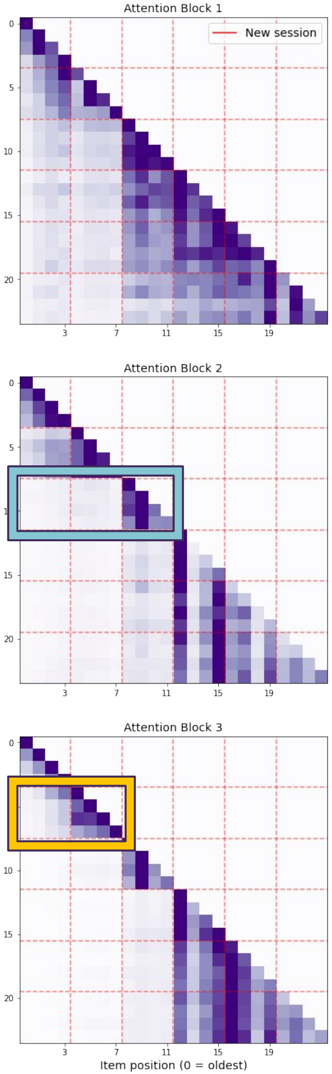

Artificial Example

- No link between sleeper sofas and jute sofas, weak link between bestselling and convertible sofas
- However, the sofas within each row of 4 are linked to each other

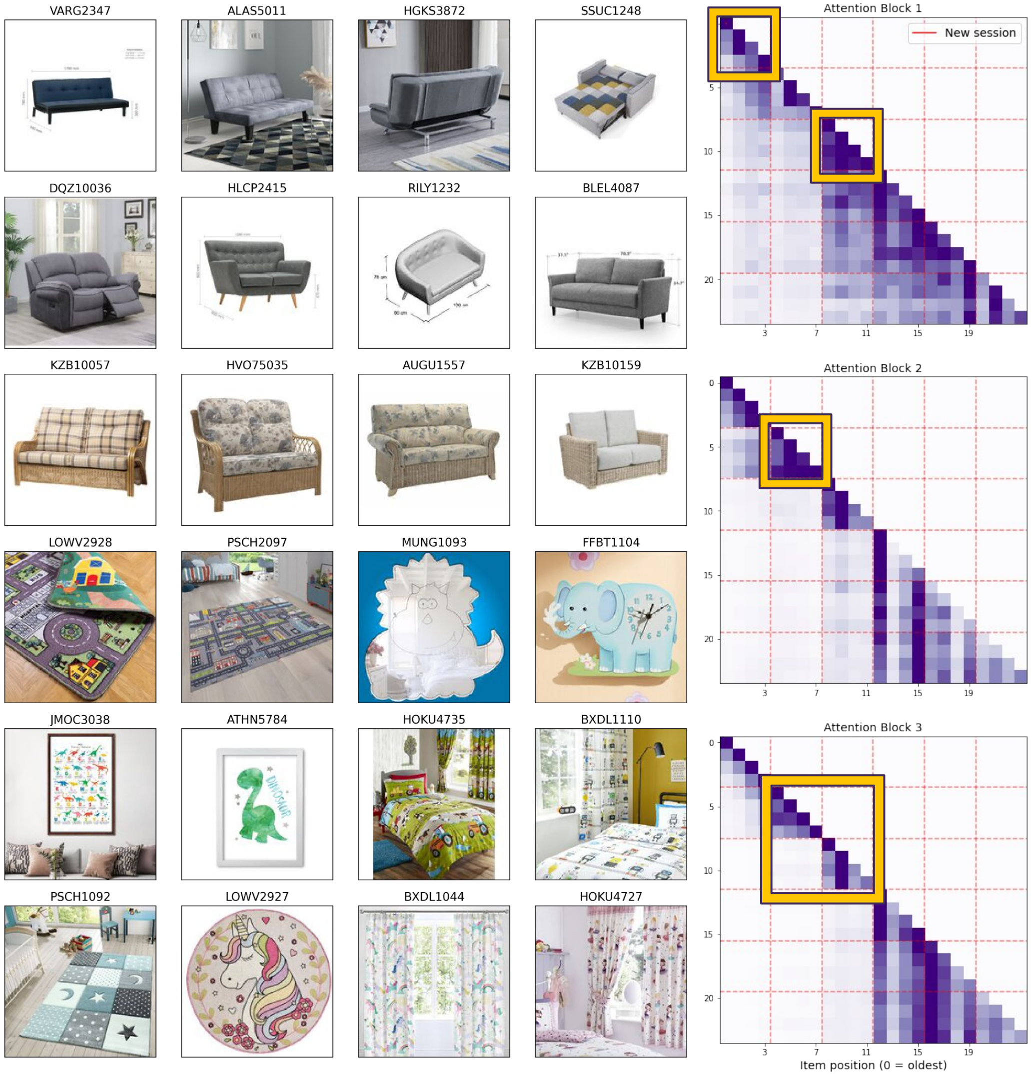

Artificial Example

- No link between sleeper sofas and jute sofas, weak link between bestselling and convertible sofas

- However, the sofas within each row of 4 are linked to each other

- Overall link between all children’s furniture even though they span different categories (rugs, mirrors, wall clocks, wall art, sheets, curtains)

⇒ So browse is not segregated by **category** but also by **style/functionality**

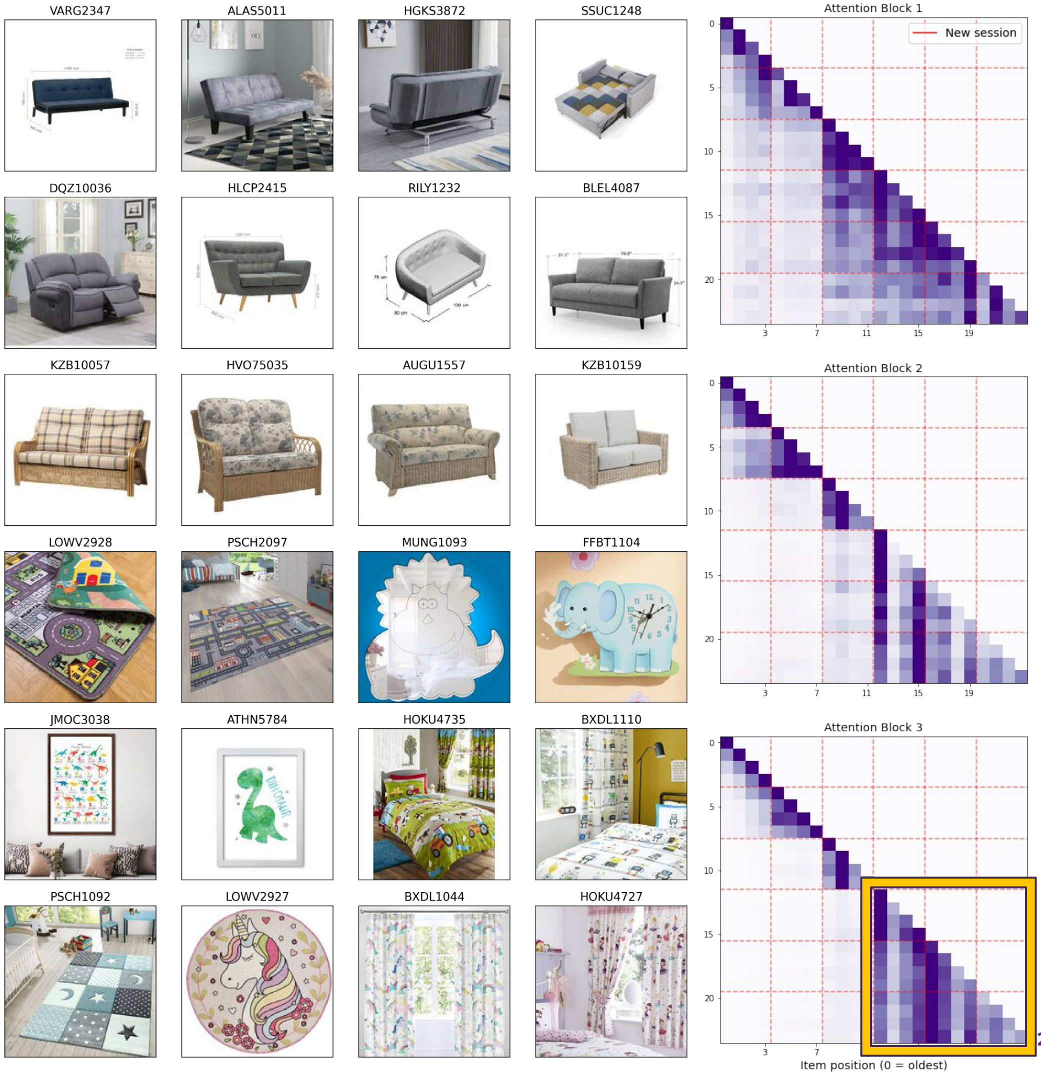

Artificial Example

- This connection between children’s furniture can persist (although weaker) even if separated by irrelevant items

- The most **relevant** context is not necessarily the most **recent** context

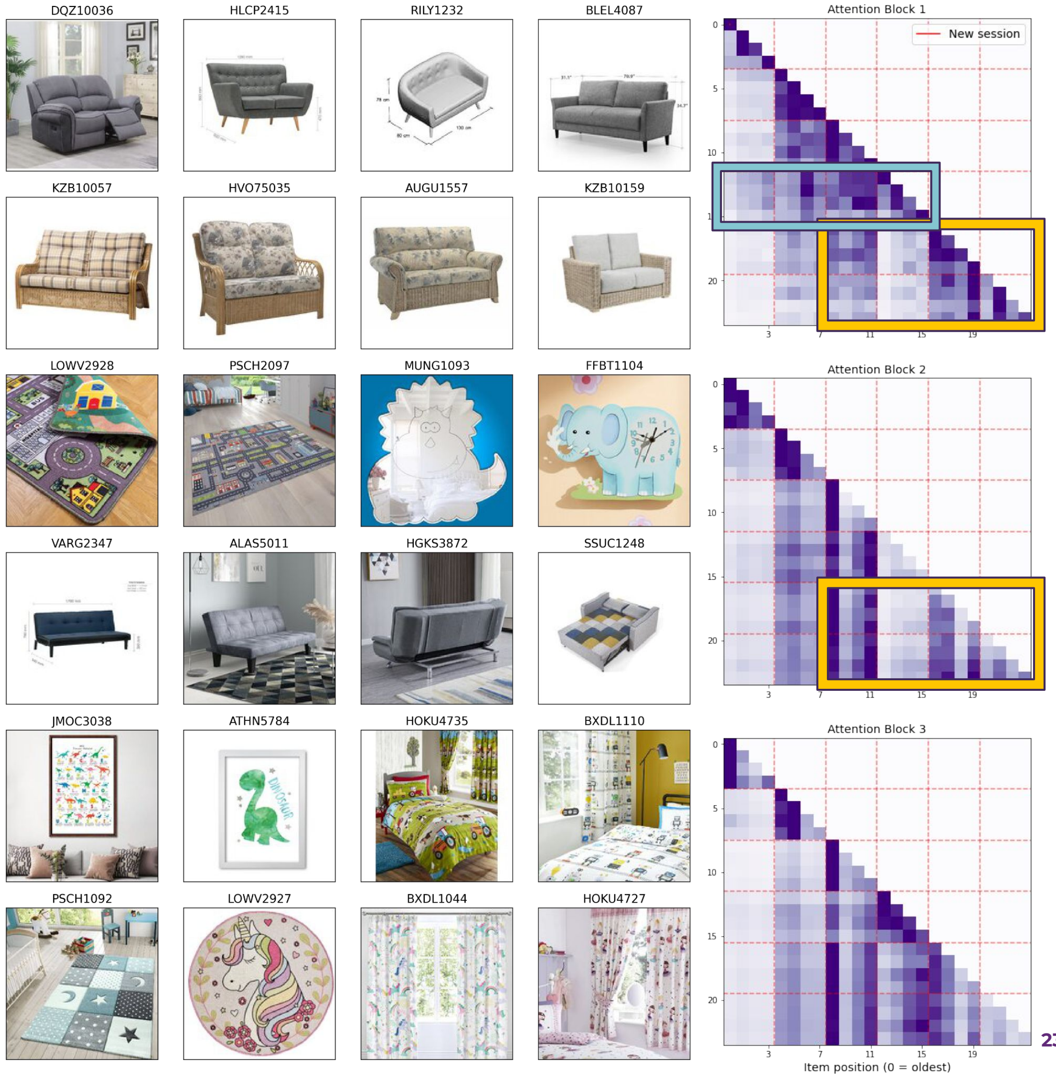

Artificial Example

- Interestingly, if you put the stereotypically-female rugs/drapes first, these are not treated as relevant to the stereotypically-male rugs/drapes, **even though they are from the same categories** (rugs/drapes)

- Still a strong link between the rugs with cars / sheets with farm equipment + wall art of dinosaurs.

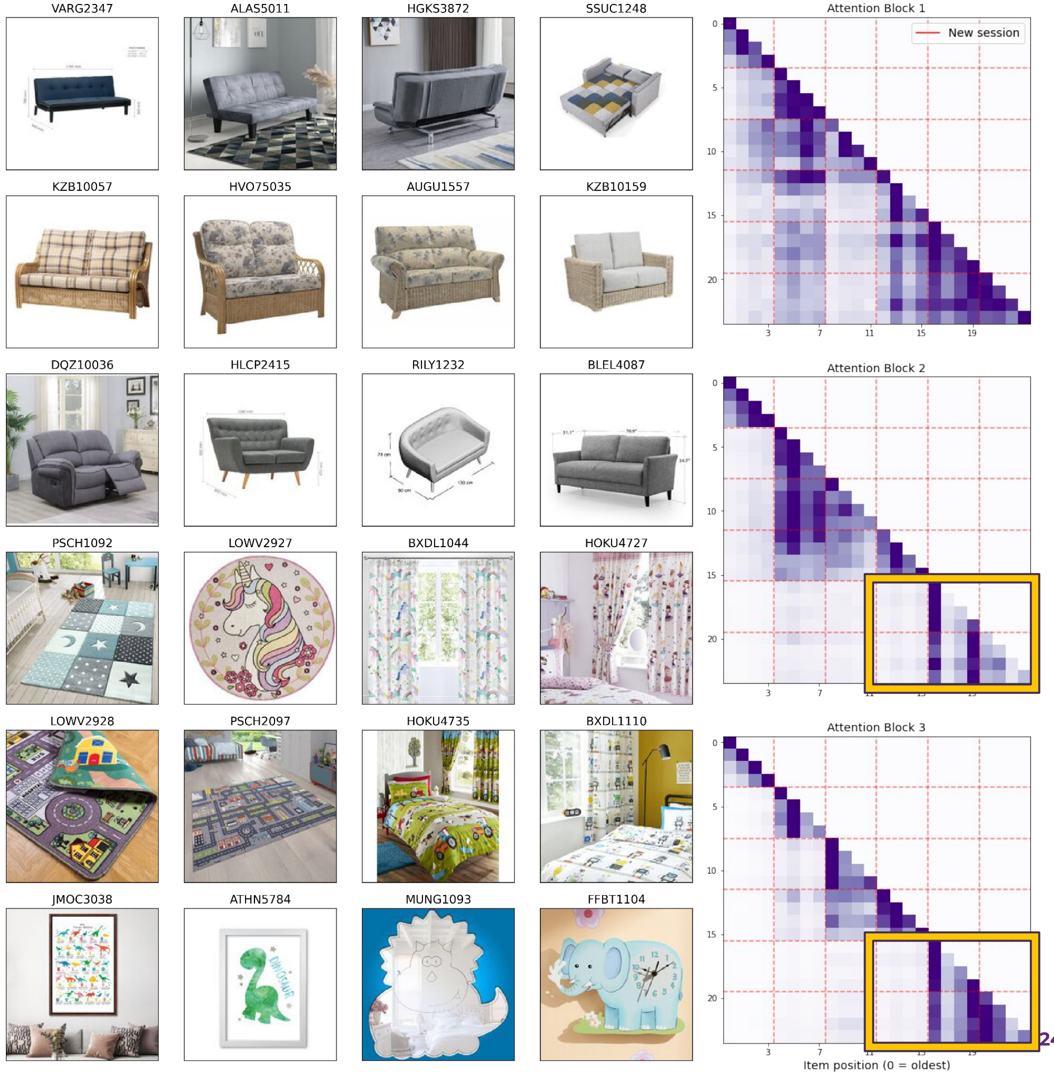

2b.

Item embeddings

Learnable attributes

- Price, Popularity signals are all learned **implicitly** by MARS
- MARS also learns category-specific attributes like table shape
    - These attributes exist on a continuum - e.g. "square" is a subset of "rectangle"; "oval" is between "round" and "rectangular"
- Because they are implicitly learned, these attributes can be all learned at the same time, and MARS learns to weight each of these independently

Item Embeddings for Coffee Tables

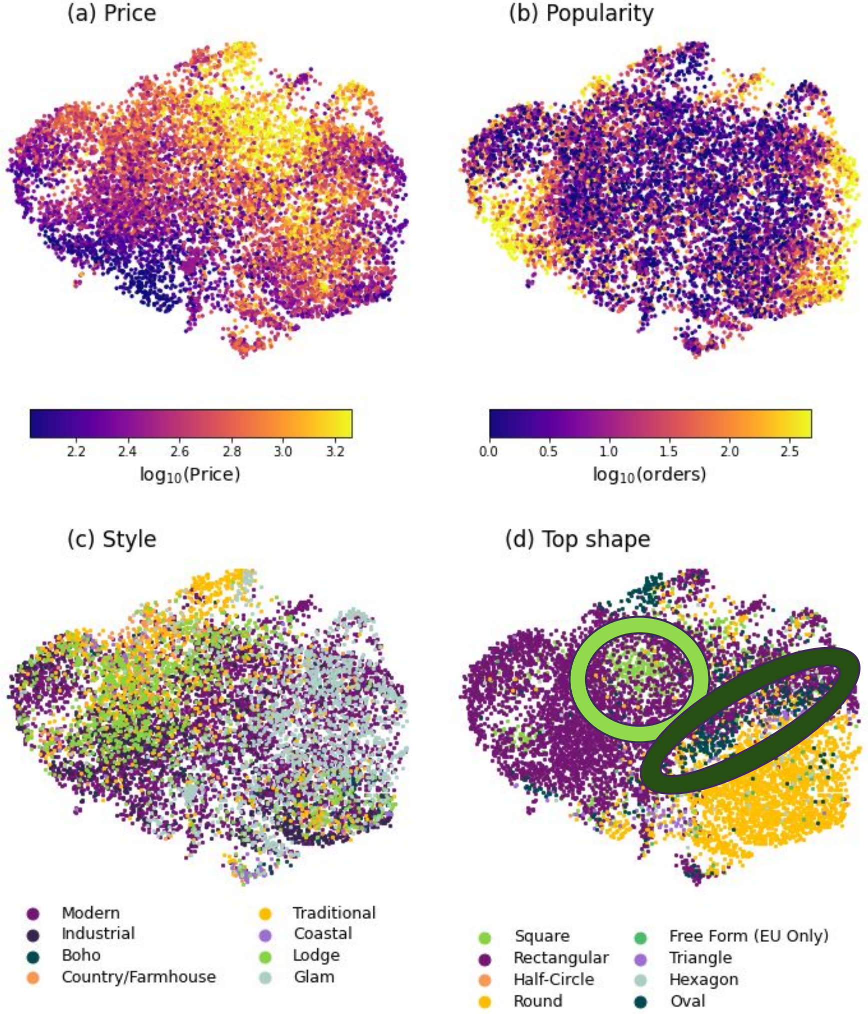

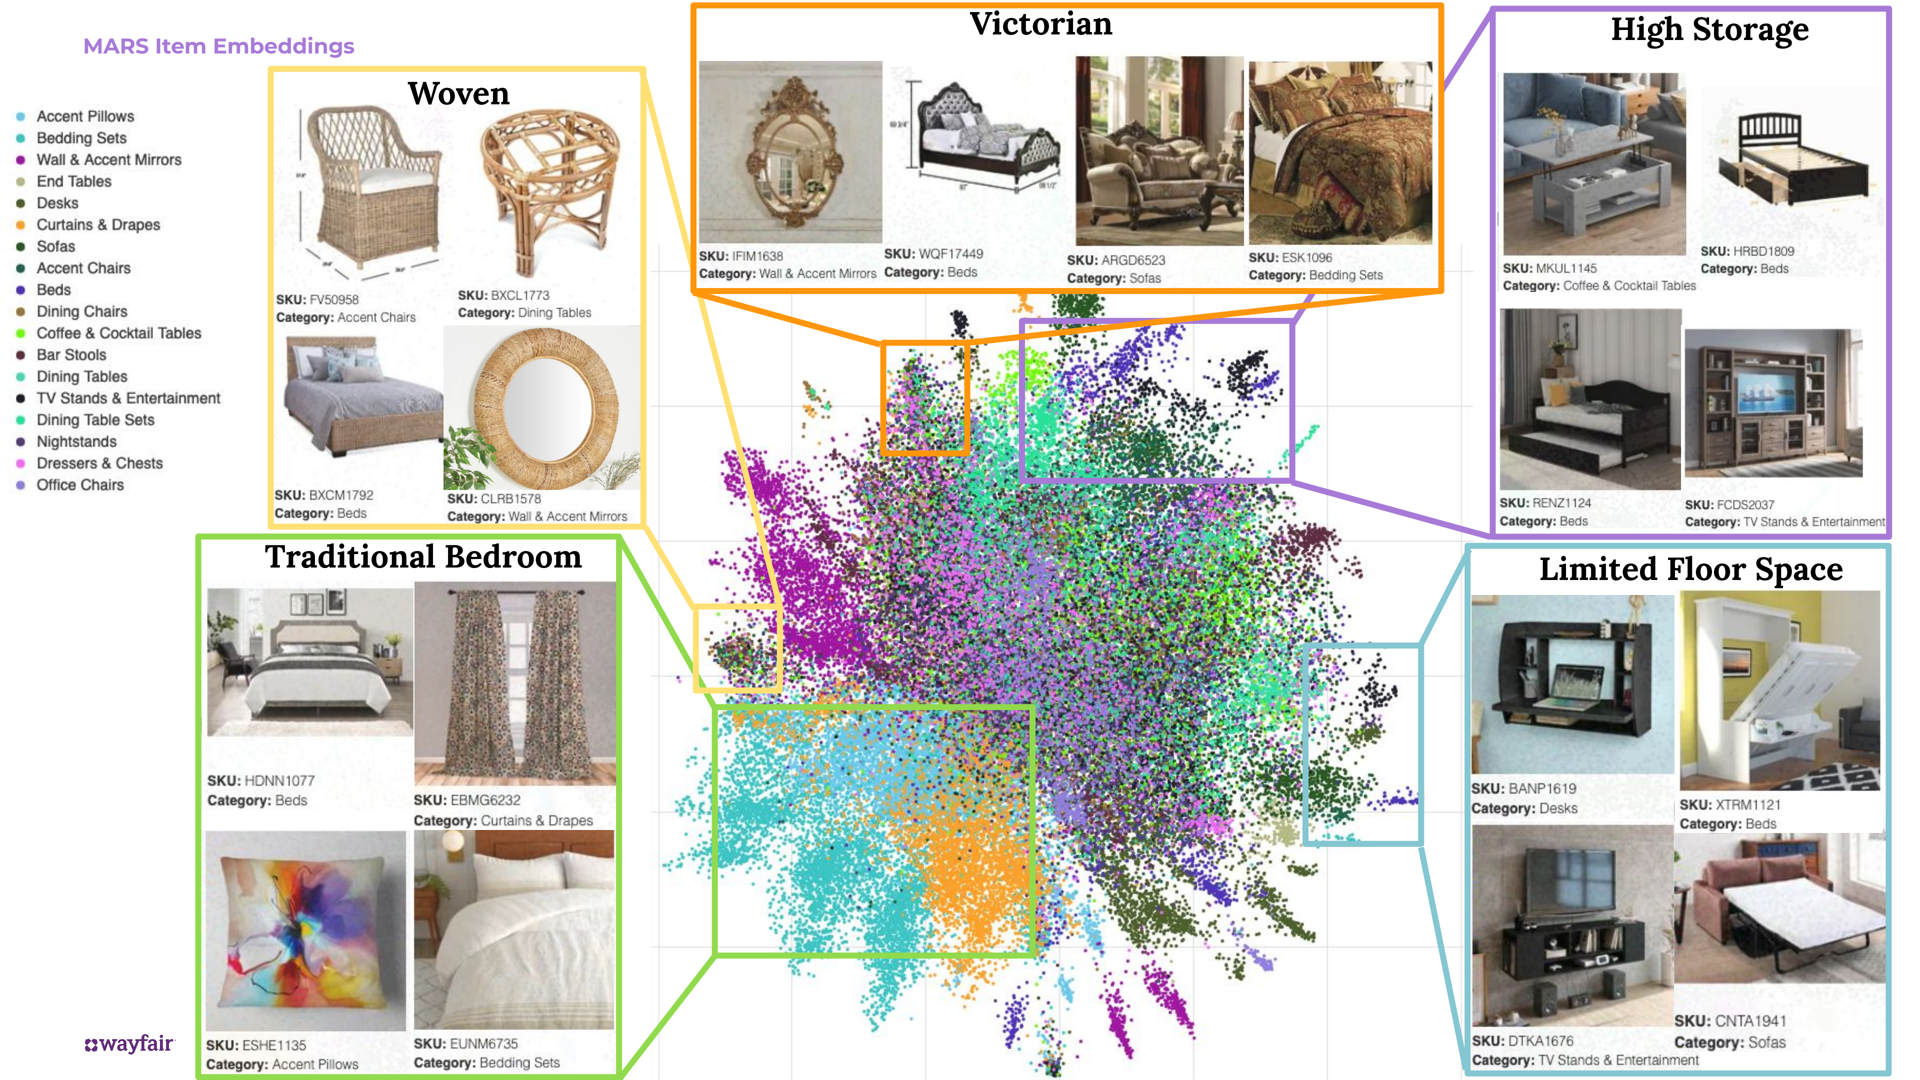
MARS Item Embeddings
Accent Pillows
Bedding Sets
Wall & Accent Mirrors
End Tables
Desks
Curtains & Drapes
Sofas
Accent Chairs
Beds
Dining Chairs
Coffee & Cocktail Tables
Bar Stools
Dining Tables
TV Stands & Entertainment
Dining Table Sets
Nightstands
Dressers & Chests
Office Chairs
Woven
SKU: FV50958
Category: Accent Chairs
SKU: BXCL1773
Category: Dining Tables
SKU: BXCM1792
Category: Beds
SKU: CLRB1578
Category: Wall & Accent Mirrors
Victorian
SKU: IFIM1638
Category: Wall & Accent Mirrors
SKU: WQF17449
Category: Beds
SKU: ARGD6523
Category: Sofas
SKU: ESK1096
Category: Bedding Sets
High Storage
SKU: MKUL1145
Category: Coffee & Cocktail Tables
SKU: HRBD1809
Category: Beds
SKU: RENZ1124
Category: Beds
SKU: FCDS2037
Category: TV Stands & Entertainment
Traditional Bedroom
SKU: HDNN1077
Category: Beds
SKU: EBMG6232
Category: Curtains & Drapes
SKU: ESHE1135
Category: Accent Pillows
SKU: EUNM6735
Category: Bedding Sets
Limited Floor Space
SKU: BANP1619
Category: Desks
SKU: XTRM1121
Category: Beds
SKU: DTKA1676
Category: TV Stands & Entertainment
SKU: CNTA1941
Category: Sofas
wayfair

2c.

Putting item and position embeddings together - adapting to changing customer preferences

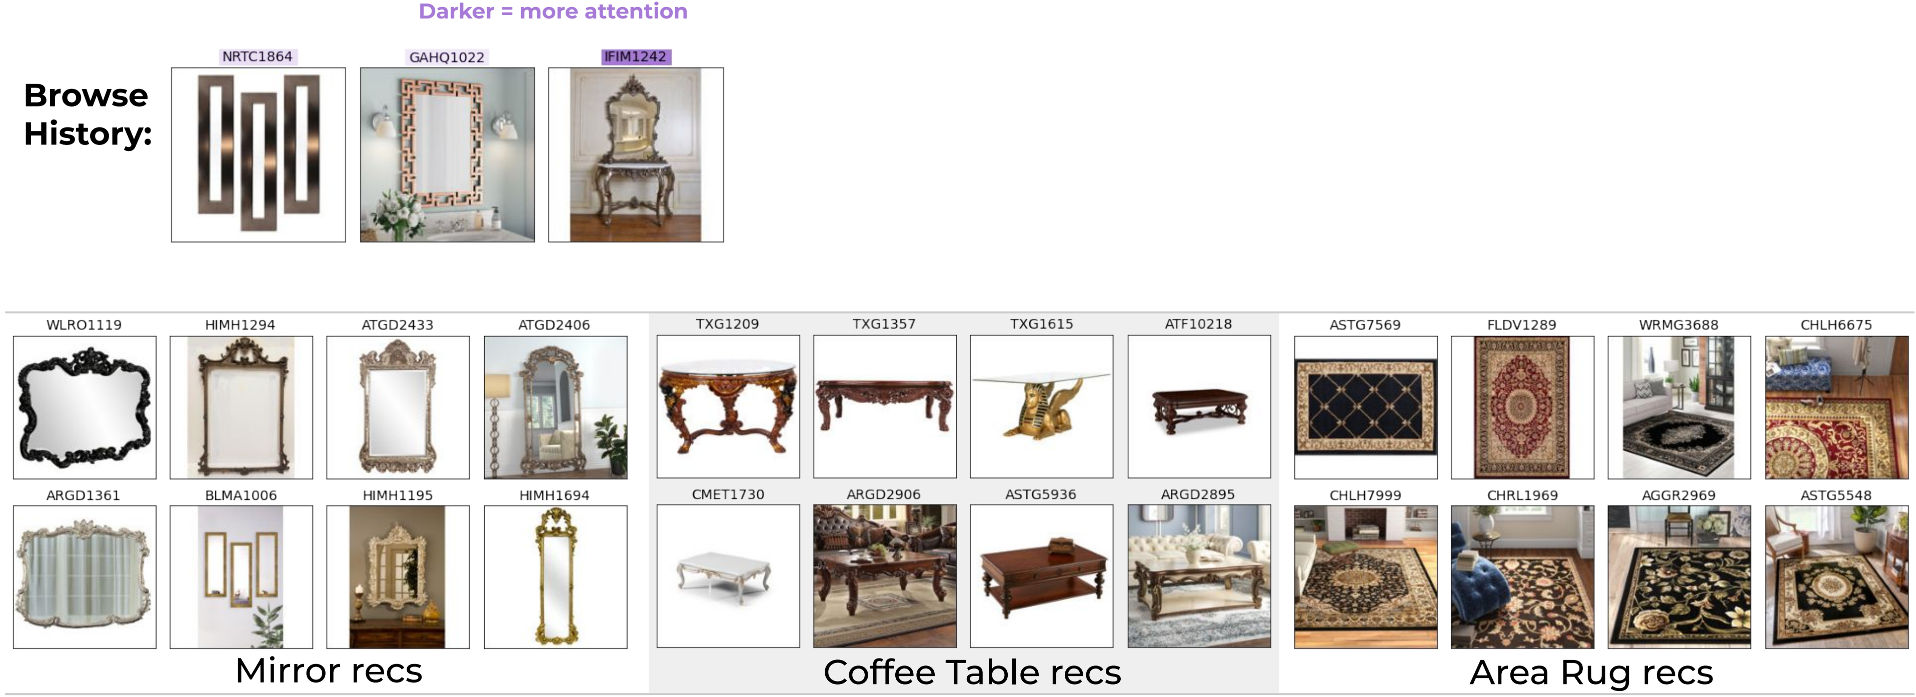
Darker = more attention
Browse History:
NRTC1864
GAHQ1022
IFIM1242
WLRO1119
HIMH1294
ATGD2433
ATGD2406
ARGD1361
BLMA1006
HIMH1195
HIMH1694
Mirror recs
TXG1209
TXG1357
TXG1615
ATF10218
CMET1730
ARGD2906
ASTG5936
ARGD2895
Coffee Table recs
ASTG7569
FLDV1289
WRMG3688
CHLH6675
CHLH7999
CHRL1969
AGGR2969
ASTG5548
Area Rug recs

Browse History:

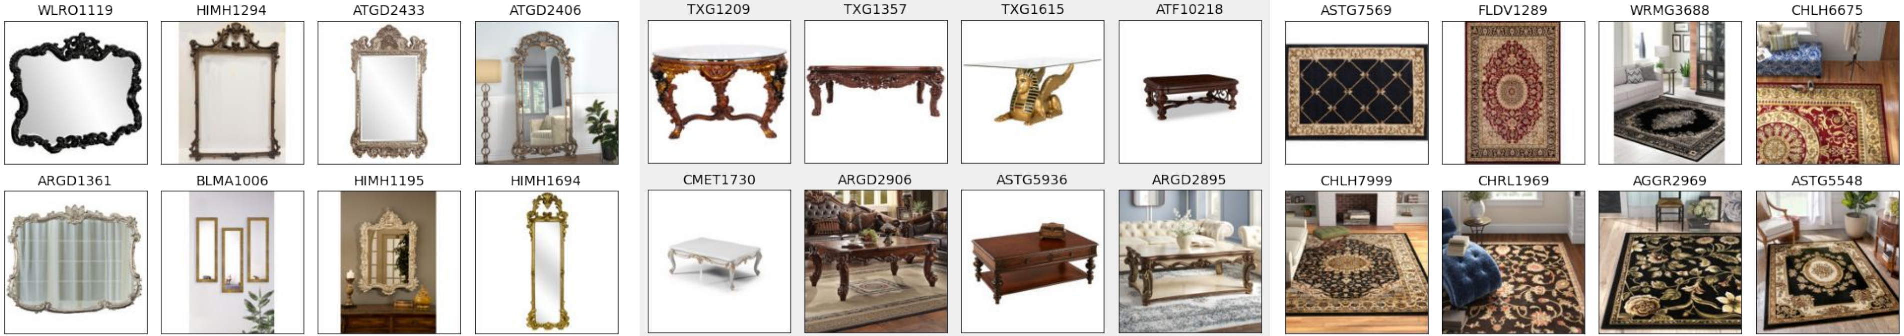

Mirror recs

Coffee Table recs

Area Rug recs

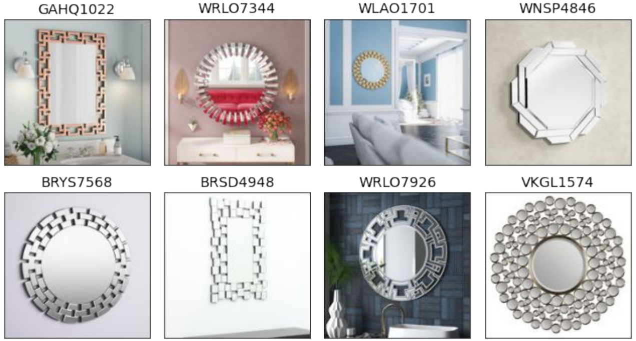

Compared to matrix factorization (non-sequential)

- Wrongly uses majority browse for mirror recs

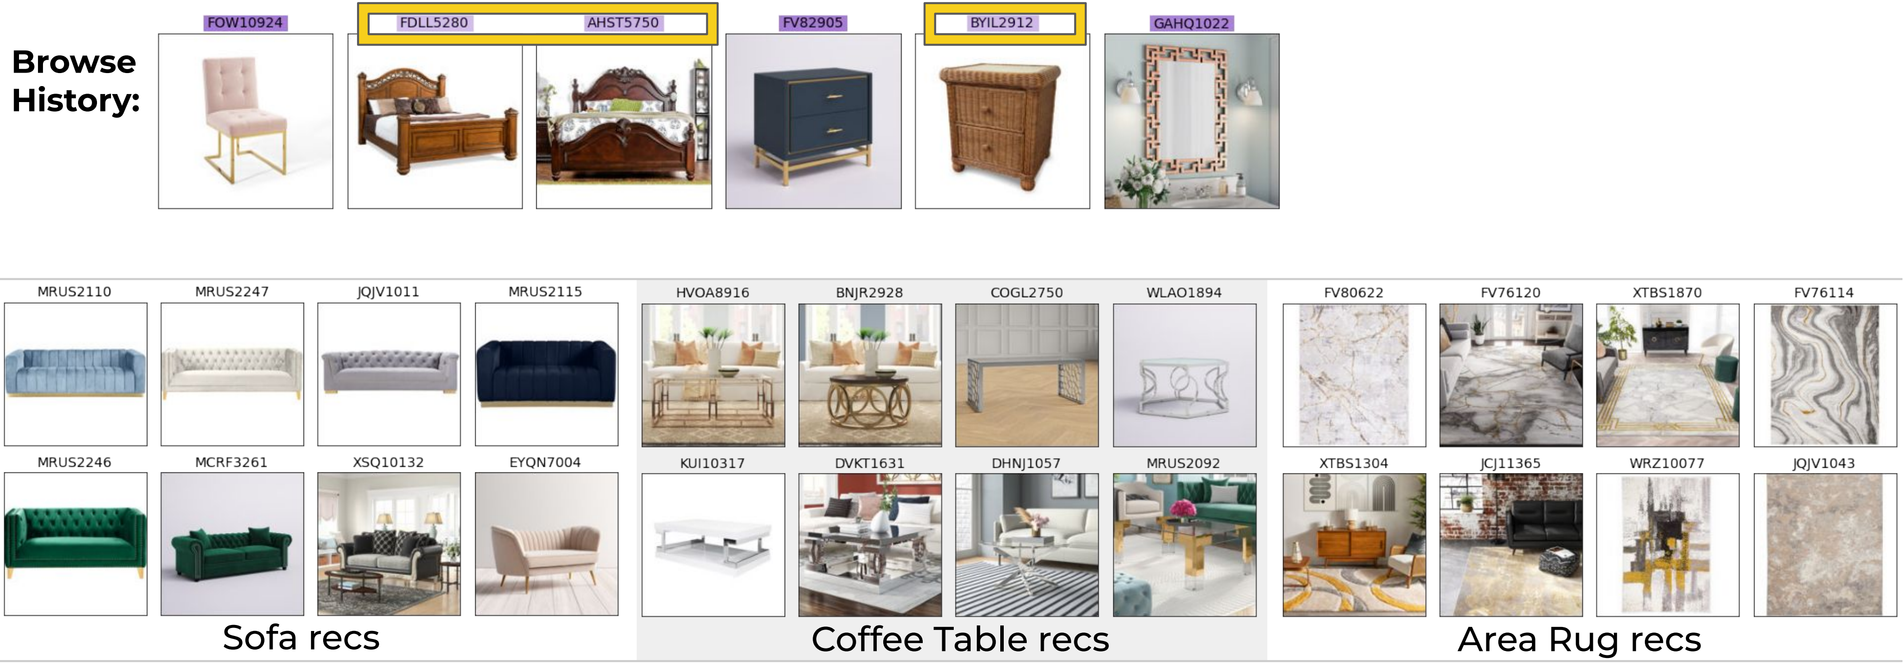

Notes:

- MARS can learn to ignore irrelevant browse (items 2, 3, 5)

<u>Notes</u>:

- MARS can learn to ignore irrelevant browse (items 2, 3, 5)
- Has learned the **material** (velvet) from item 1; **leg style** (metallic; both ‘sled base’ and straight’ legs) from item 1 and 4

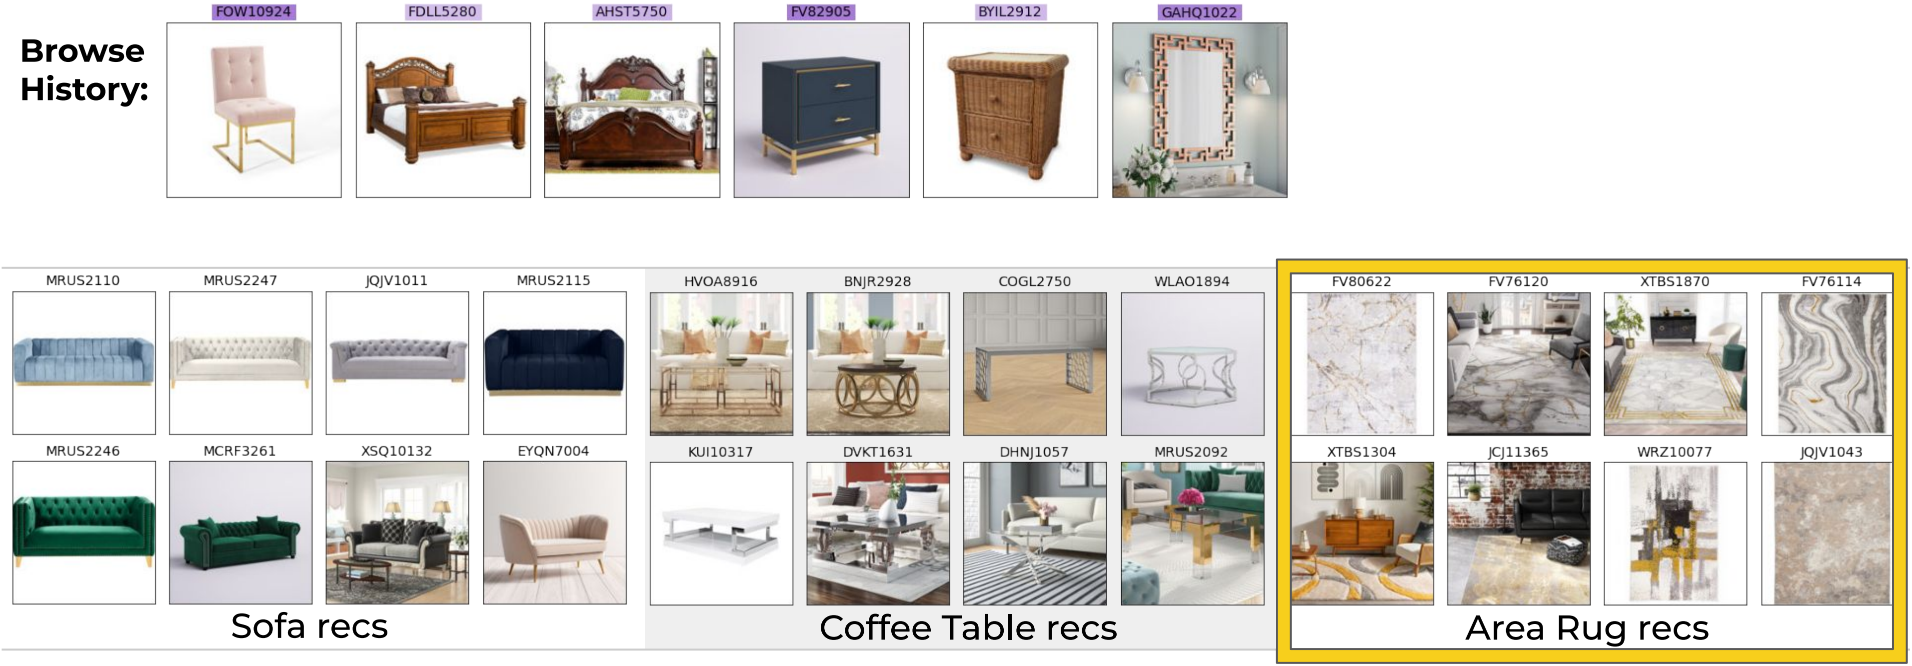

<u>Notes</u>:
- MARS can learn to ignore irrelevant browse (items 2, 3, 5)
- Has learned the **material** (velvet) from item 1; **leg style** (metallic; both ‘sled base’ and straight’ legs) from item 1 and 4
- Even the area rug recommendations have metallic colors (gold/silver)

3.

User Intent

Resurfacing Items

Users typically like to view the same item

- A real customer's browsing history; ordered items circled
- A customer's actual order is
  - not necessarily most viewed
  - not necessarily last-viewed
- So our model should learn to resurface, but how/when?

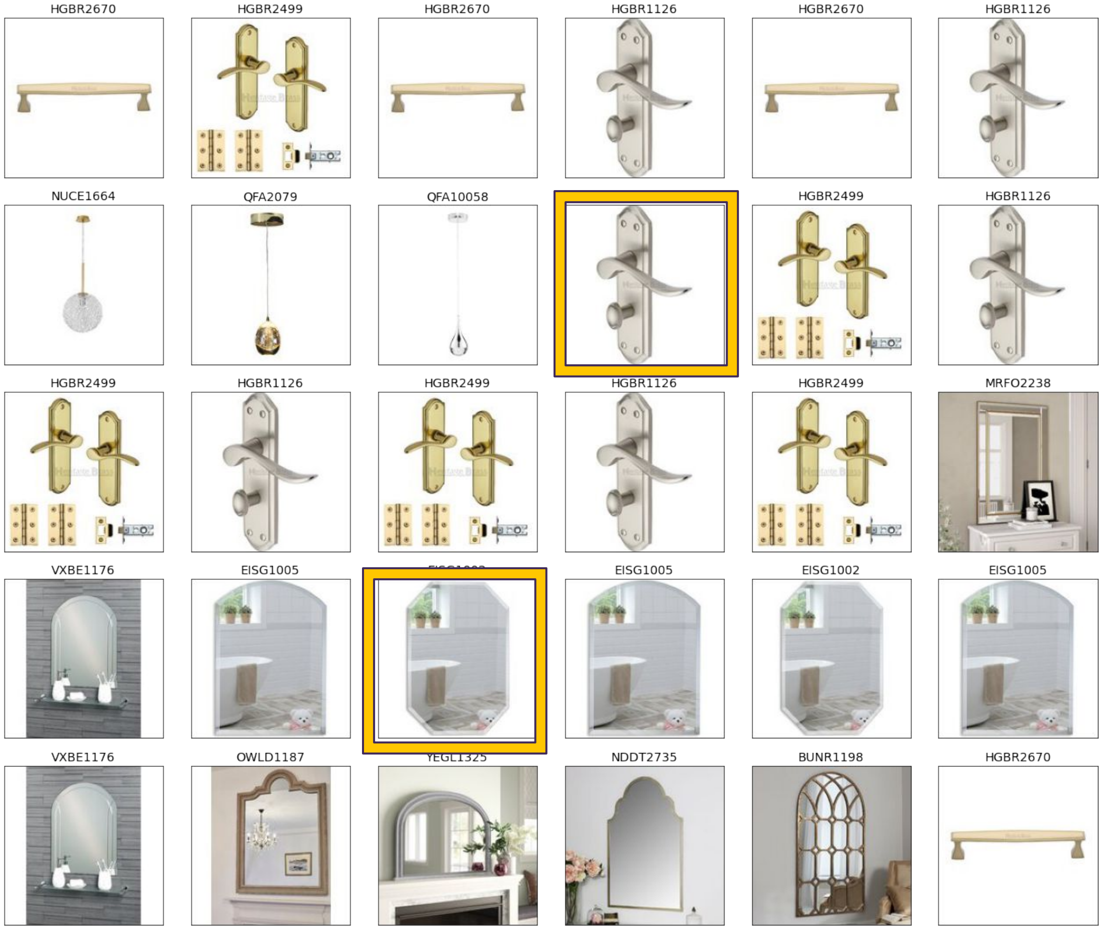

Injecting resurfacing awareness

What item should we actually resurface?

- Not necessarily most-viewed
- Not necessarily last-viewed
- But rather most likely to be ordered **given it is viewed**

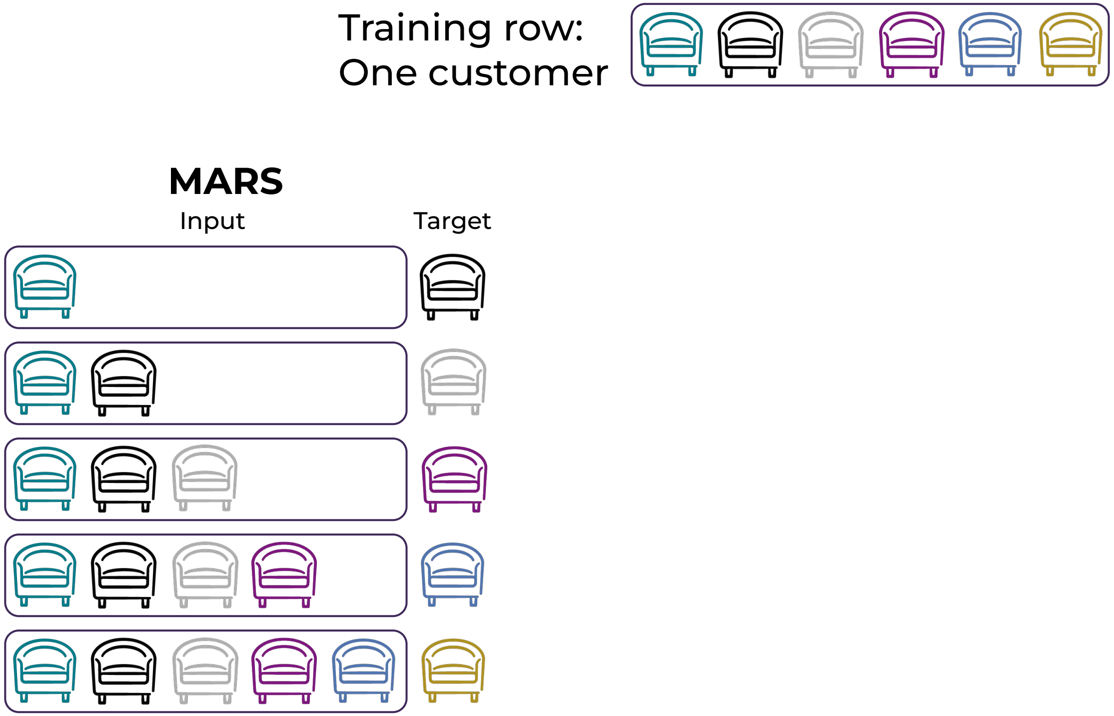

Injecting resurfacing awareness

What item should we actually resurface?

- Not necessarily most-viewed
- Not necessarily last-viewed
- But rather most likely to be ordered **given it is viewed**

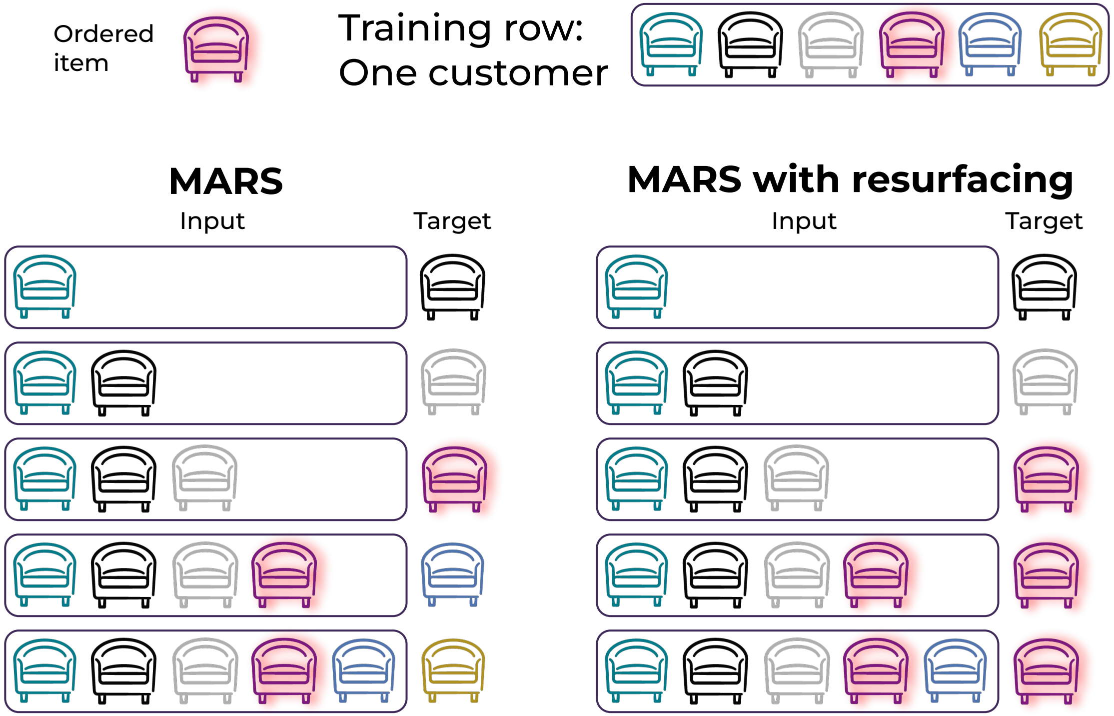

Results

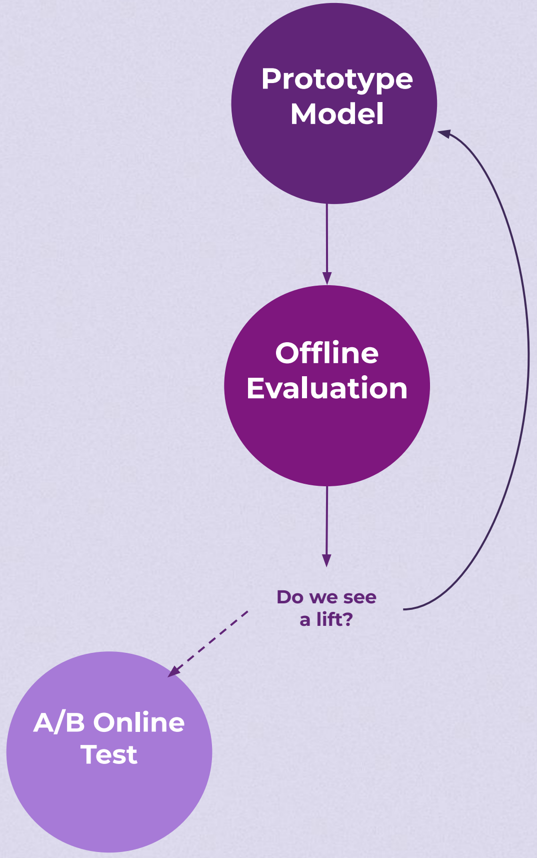

- Using historic orders
- Can use any metric - we use MRR as a proxy for CVR
- Typically 10k-100k customers
- Slightly biased towards control

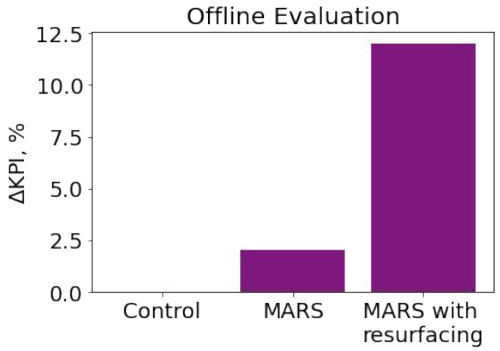

- Split 50/50 of real-time traffic into control/variation

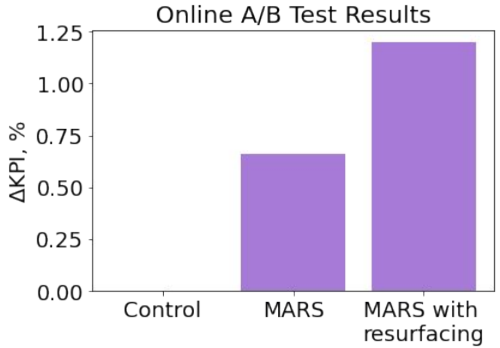

Takeaways

Sequential information is (rich) context

Dynamic but powerful, there is a lot of style/price signal hidden in the browse history, which we can use to **transfer style/functionality/material** signals across recs for different furniture types

Transformers are powerful

By extracting and visualizing the **self-attention** weights, we can intuit good guesses (always with a grain of salt!) about how the model is learning / what parts of the **browse context** the model considers to be actually **relevant** (vs. what it can ignore)

Develop an understanding of the customer and tweak models

How does a customer get from **view** to **order**? How can we incorporate this info into our models (e.g. customers like to buy recently-viewed items)?

- Our Tech Blog
- Careers