3. DEMONSTRATION & USE CASES

In this section, we demonstrate different use-cases and real-world scenarios where our system might be used to help a user to explore the visual space. Each of these use-cases can be naturally handled by *Fashionista*, and to the best of our knowledge can not be handled by existing systems.

3.1 Querying for fashionability advice

Rachel is hesitant to purchase a dress on *Amazon*. She personally likes it but is not sure whether it is consistent with the current 'fashion zeitgeist,' i.e., its fashionability. She searches for the dress on *Fashionista*. *Fashionista* demonstrates the fashionability evolution of the dress's appearance in the past decade, and Rachel finds that it's gaining popularity in recent years. Armed with this information, she decides to purchase the dress immediately.

3.2 Searching substitutes with similar styles

Lucy finds a pair of shoes (e.g. *B004V7858C*) on *Amazon*, but wants to compare them against alternatives that are similar in appearance, but potentially have a preferable brand, price, or rating (etc.). Among items that are 'frequently bought together' (as recommended by *Amazon*), Lucy finds they tend to vary too much in appearance, or otherwise are not visually attractive. Using *Fashionista*, she quickly retrieves hundreds of visually similar items and finds a shoe of the same style but with a preferable brand.

3.3 Finding complements for outfit generation

Angelina is herself a fashionista who just bought a beautiful t-shirt on *Amazon*, and she wants to find some pants and shoes that can go together with it, that is, to generate an outfit with a consistent visual style. She uses *Fashionista*'s category filter to limit the search results to be within the two categories under consideration. According to the query conditions, the system retrieves the nearest neighbors of the t-shirt, according to the learned visual 'style space,' each represented with different colors denoting their visual popularity. Angelina explores these items and finds several with high fashionability scores that match the style of the t-shirt.

3.4 I'm feeling fashionable!

Unlike the above scenarios, Christina doesn't know what she really wants, but just intends to explore some fashionable items. She clicks the 'I'm feeling fashionable!' button on *Fashionista*, and the system returns a randomly selected hotspot in the visual space that currently has a high popularity (fashionability). She decides to focus on fashionable coats, using the provided category filter. Some attractive coats are surfaced, and Christina decides to buy one of them using the redirection links.

3.5 Querying for statistical fashion trends

Jennifer is a fashion designer who cares about the trends of contemporary fashion. Using *Fashionista*, she quickly checks the distribution of popular appearances over hundreds of thousands of clothing and accessory items. She zooms into some areas that interest her the most and observes the corresponding fashion trend evolution during the past decade. She identifies certain trends that are likely to gain popularity in the near future and decides to design products that fit them.

4. RELATED WORK

Most relevant to our work are existing e-commerce systems like *Amazon*, *eBay*,[4] *Walmart*,[5] etc. In spite of the wide successes achieved by such systems, they are not designed for exploration of visually relevant items. Although they are able to recommend items that are frequently bought or viewed together, they are unable to tease-apart the underlying reasons including functionality, complementarity, visual compatibility, and so forth. This limitation makes it difficult for these systems to address the task of recommending visually consistent items, nor do they provide an interface for efficient exploration within a visual space as we do in this work.

There are also image search engines like *Google Images* that are either based on keywords or image similarity comparison. Such systems differ from ours in two key ways. First, such systems are general-purpose and thus are not suitable for the specific applications considered here. Second, our model maps items into a low-dimensional visual space such that nearby items have been evaluated similarly (by users) in terms of their appearance. This improves upon traditional image similarity comparison methods by performing comparisons only on those dimensions that are discovered to be relevant to people's decisions.

5. CONCLUSION AND FUTURE WORK

In this paper, we built a demo system, *Fashionista*, with a user-friendly graphical interface to search and explore visually similar and fashionable items. Experimentally, we found that *Fashionista* can index and search large-scale real-world corpora efficiently and effectively. In the future, we hope to enable our system to automatically generate outfits according to *personalized* visual preferences and fashion trends. This will be feasible once we are able to observe users' interactions with *Fashionista*.

6. REFERENCES

[1] J. J. McAuley, C. Targett, Q. Shi, and A. van den Hengel, "Image-based recommendations on styles and substitutes," in *SIGIR*, 2015.

[2] R. He and J. McAuley, "Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering," in *WWW*, 2016.

[3] R. He and J. McAuley, "VBPR: visual bayesian personalized ranking from implicit feedback," in *AAAI*, 2016.

[4] C. Lin, J. Lu, T. W. Ling, and B. Cautis, "Lotusx: a position-aware xml graphical search system with auto-completion," in *ICDE*, 2012.

[5] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *NIPS*, 2012.

[6] R. Pan, Y. Zhou, B. Cao, N. N. Liu, R. Lukose, M. Scholz, and Q. Yang, "One-class collaborative filtering," in *ICDM*, 2008.

[7] L. van der Maaten, "Accelerating t-SNE using tree-based algorithms," *JMLR*, 2014.

[4]http://www.ebay.com/

[5]http://www.walmart.com/