

From Paris to Berlin: Discovering Fashion Style Influences Around the World

Ziad Al-Halah¹

ziadhlh@gmail.com

¹The University of Texas at Austin

Kristen Grauman^{1,2}

grauman@cs.utexas.edu

²Facebook AI Research

Abstract

The evolution of clothing styles and their migration across the world is intriguing, yet difficult to describe quantitatively. We propose to discover and quantify fashion influences from everyday images of people wearing clothes. We introduce an approach that detects which cities influence which other cities in terms of propagating their styles. We then leverage the discovered influence patterns to inform a forecasting model that predicts the popularity of any given style at any given city into the future. Demonstrating our idea with GeoStyle—a large-scale dataset of 7.7M images covering 44 major world cities—we present the discovered influence relationships, revealing how cities exert and receive fashion influence for an array of 50 observed visual styles. Furthermore, the proposed forecasting model achieves state-of-the-art results for a challenging style forecasting task, showing the advantage of grounding visual style evolution both spatially and temporally.

1. Introduction

“The influence of Paris, for instance, is now minimal. Yet a lot is written about Paris fashion.”—Geoffrey Beene

The clothes people wear are a function of personal factors like comfort, taste, and occasion—but also wider and more subtle influences from the world around them, like changing social norms, art, the political climate, celebrities and style icons, the weather, and the mood of the city in which they live. Fashion itself is an evolving phenomenon because of these changing influences. What gets worn continues to change, in ways fast, slow, and sometimes cyclical.

Pinpointing the influences in fashion, however, is non-trivial. To what extent did the runway styles in Paris last year affect what U.S. consumers wore this year? How much did the designs by J. Crew influence those created six months later by Everlane, and vice versa? How long does it take for certain trends favored in New York City to migrate to Austin, if they do at all? And how did the infamous cerulean sweater worn by the protagonist in the movie *The*

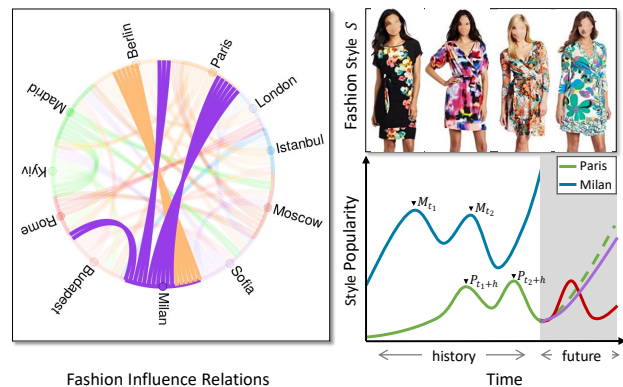


Figure 1: Styles propagate according to certain patterns of influence around the world. For example, the trajectory of a given style’s popularity in Milan may foreshadow its trajectory in Paris some months later. Our idea is to discover style influence relations worldwide (left) and leverage them to accurately forecast future trends per location (right). Whereas forecasting without regard to geographic influence can falter in the presence of complex trends (red curve), using discovered influence information (e.g., *Milan influences Paris* in fashion style *S*) yields better forecasts (purple curve).

Devil Wears Prada make its way into her closet?¹

To quantitatively answer such questions would be valuable to both social science and business forecasts, yet it remains challenging. Clothing sales records or social media “likes” offer some signal about how tastes are shifting, but they are indirect and do not reveal the sources of influence.

We contend that images are exactly the right data to answer such questions. Unlike vendors’ purchase data, other non-visual metadata, or hype from haute couture designers, everyday photos of what people are wearing in their daily life provide a unfiltered glimpse of current clothing styles “on the ground”. Our idea is to discover fashion influence patterns in community photo collections (e.g., Instagram, Flickr), and leverage those influence patterns to forecast future style trends conditioned on the place in the world.

¹The Devil Wears Prada: Cerulean <https://bit.ly/3dBAQ5W>

While fashion influences exist along several axes, we focus on worldwide geography to capture spatio-temporal influences. Specifically, we aim to discover *which cities influence which other cities* in terms of propagating their clothing styles, and *with what time delay*.

To this end, we introduce an approach to discover geographical style influences from photos. First, we extract a vocabulary of visual styles from unlabeled, geolocated and timestamped photos of people. Each style is a mixture of detected visual attributes. For example, one style may capture short floral dresses in bright colors (Fig. 1) while another style captures preppy collared shirts. Next, we record the past trajectories of the popularity of each style, meaning the frequency with which it is seen in the photos over time. Then, we identify two key properties of an influencer—time precedence and novelty—and use a statistical measure that captures these properties to calculate the degree of influence between cities. Next, we introduce a neural forecasting model that exploits the influence relationships discovered from photos to better anticipate future popular styles in any given location. Finally, we propose a novel coherence loss to train our model to reconcile the local predictions with the global trend of a style for consistent forecasts. We demonstrate our approach on the large-scale GeoStyle dataset [28] comprised of everyday photos of people with a wide coverage of geographic locations.

Our results shed light on the spatio-temporal migration of fashion trends across the world—revealing which cities are exerting and receiving more influence on others, which most affect global trends, which contribute to the prominence of a given style, and how a city’s degree of influence has itself changed over time. Our findings hint at how computer vision can help democratize our understanding of fashion influence, sometimes challenging common perceptions about what parts of the world are driving fashion (consistent with designer Geoffrey Beene’s quote above).

In addition, we demonstrate that by incorporating influence, the proposed forecasting model yields state-of-the-art accuracy for predicting the future popularity of styles. Unlike prior work that learns trends with a monolithic worldwide model [1] or independent per city models [28], our geo-spatially grounded predictions catch the temporal dependencies between when different cities will see a style climb or dip, producing more accurate forecasts (see Fig. 1).

2. Related Work

Visual fashion analysis, with its challenging vision problems and direct impact on our social and financial life, presents an attractive domain for vision research. In recent years, many aspects of fashion have been addressed in the computer vision literature, ranging from learning fashion attributes [2, 3, 6, 7, 26], landmark detection [39, 41], cross-domain fashion retrieval [19, 10, 42, 23], body shape

and size based fashion suggestions [31, 14, 17], virtual try-on [38, 8], clothing recommendation [25, 30, 40, 18], inferring social cues from people’s clothes [35, 32, 24], outfit compatibility [16, 12], visual brand analysis [22, 11], and discovering fashion styles [21, 36, 1, 15]. Our work opens a new avenue for visual fashion understanding: modeling influence relations in fashion directly from images.

Statistics of styles Analyzing styles’ popularity in the past gives a window on people’s preferences in fashion. Prior work considers how the frequency of attributes (*e.g.*, floral, neon) changed over time [37], and how trends in (non-visual) clothing meta-data changed for the two cities Manila and Los Angeles [34]. Qualitative studies suggest how collaborative filtering recommendation models can account for past temporal changes of fashion [13] or what cities exhibit strong style similarities [20]. However, all this prior work analyzes style popularity in an “after the fact” manner, and looks only qualitatively at past changes in style trends. We propose to go beyond this historical perspective to forecast future changes in styles’ popularity along with supporting quantitative evaluation.

Trend forecasting Only limited prior work explores forecasting visual styles into the future [1, 28]. The FashionForward model [1] uses fashion styles learned from Amazon product images to train an exponential smoothing model for forecasting, treating the products’ transaction history (purchases) as a proxy for style popularity. Similarly, the GeoStyle project [28] uses a seasonal forecasting model to predict changes in style trends per city and highlight unusual events. Both prior models assume that style trends in different cities are independent from one another and can be modeled monolithically [1] or in isolation [28]. In contrast, we introduce a novel model that accounts for influence patterns discovered across different cities. Our concept of fashion influence discovery is new, and our resulting forecasting model outperforms state-of-the-art.

Influence modeling To our knowledge, no previous work tackles influence modeling in the visual fashion domain. The closest study looks at the correlation among attributes popular in New York fashion shows and those attributes seen in street photos, as a surrogate for fashion shows’ impact on people’s clothing; however, no influence or forecasting model is developed [7]. Outside the fashion domain, models for influence are developed for connecting text in news articles [33], linking video subshots for summarization [27], or analyzing intellectual links between major AI conferences from their papers [5]. Our method is the first to model influence in visual fashion trends. We propose an influence model that is grounded by forecasting accuracy. Through our evaluation, we show that our model discovers interesting influence patterns in fashion that go beyond simple correlations, and we analyze influence on multiple axes to discover locally and globally influential players.

3. Visual Style Influence Model

We propose an approach to model influence in the visual fashion domain. Starting with images of fashion garments, 1) we learn a visual style model that captures the fine-grained properties common among the garments; then 2) we construct style popularity trajectories by leveraging images’ temporal and spatial meta information; 3) we model the influence relations between different locations (cities) for a given visual style; Finally, 4) we introduce a forecasting model that utilizes the learned influence relations together with a coherence loss for consistent and accurate predictions of future changes in style popularity.

3.1. Visual Fashion Styles

Our model captures the fashion influence among different locations in the world. We begin by discovering a set of visual fashion *styles* from images of people’s garments in their everyday life. As discussed above, such photos offer an unfiltered glimpse of what people are wearing around the world. We use the 7.7M-image GeoStyle dataset [28] from Instagram and Flickr as source data.

Let $X = \{x_i\}^N$ be a set of clothing images. We first learn a semantic representation that captures the elementary fashion attributes like colors (*e.g.* cyan, green), patterns (*e.g.* stripes, dots), shape (*e.g.* vneck, sleeveless) and garment type (*e.g.* shirt, sunglasses). Given a fashion attribute model $f_a(\cdot)$ (*e.g.*, a CNN) trained on a set of disjoint labeled images, we can then represent each image in X with $\mathbf{a}_i = f_a(x_i)$, where $\mathbf{a} \in \mathbb{R}^M$ is a vector of M visual attribute probabilities.

Next, we learn a set of fashion styles $S = \{S_k\}^K$ that capture distinctive attribute combinations using a Gaussian mixture model (GMM) of K components. Fig. 2 shows a set of three fashion styles discovered by this style model from Instagram images. Hence, given an image of a new garment x_i , the style model $f_s(\cdot)$ can predict the probabilities of that garment to be from each of the learned styles $s_i = f_s(\mathbf{a}_i)$.

Style trajectories We measure the popularity of a fashion style in a certain location through the frequency of the style in the photos of the people in that location. Specifically, given the timestamps and geolocations of the photos, we first quantize them into a meaningful temporal resolution (*e.g.* weeks, months) and locations (*e.g.* cities). Then, we construct a temporal trajectory y^{ij} for each pair of style and location (S^i, C^j):

$$y_t^{ij} = \frac{1}{|C_t^j|} \sum_{x_k \in C_t^j} p(S^i|x_k), \quad (1)$$

where C_t^j is the set of images from location C^j in the time window t , $p(S^i|x_k)$ is the probability of style S^i given image x_k based on our style model $f_s(\cdot)$, and y_t^{ij} is the popularity of style S^i in location C^j during time t . Finally, by

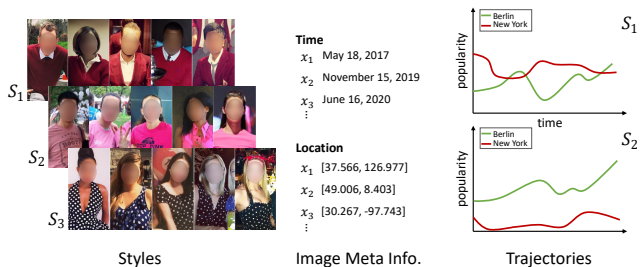


Figure 2: Style trajectories. First, we learn a set of fashion styles from everyday images (left). Then based on images’ timestamps and geolocations (center) we measure the popularity of a style at a given place (*e.g.* city) and a time period (*e.g.* week) to build up the style popularity trajectory (right).

getting all values for $t = 1, \dots, T$ we construct the temporal trajectory y^{ij} . See Fig. 2.

The GeoStyle dataset, like any Internet photo dataset, it has certain biases in terms of the demographics of the people who have uploaded photos and the locations—as discussed by the dataset creators [29]. These biases may affect the type of styles considered and their measured popularity. For example, younger generations are more likely to upload photos to Instagram and from places with easy access to high Internet bandwidth. Nonetheless, the dataset is the largest public fashion dataset with the most temporal and geographic coverage, providing a unique glimpse on people’s fashion preferences around the globe.

Next, we describe our influence model that analyzes these trajectories to discover the influence patterns among the various locations.

3.2. Influence Modeling

We propose to ground fashion influence through style popularity forecasting. This enables us to quantitatively evaluate influence using learned computational models based on real world data.

We say city C^i influences city C^j in a given fashion style S^n if our ability to accurately forecast the popularity of S^n in C^j significantly improves when taking into consideration the past popularity trend of S^n in C^i , in addition to its past popularity trend in C^j . In other words, past observations in $y_{1..t}^{ni}$ provide us with new insight on the future changes in $y_{t+1..t+h}^{nj}$ that are not available in $y_{1..t}^{nj}$.

We identify two main properties of the influencer C^i : 1) time precedence, that is the influencer city’s changes happen before the observed impact on the influenced city and 2) novelty, that is the influencer city has novel past information not observed in the history of the influenced city.

A naive approach to capture such relations is to use a multivariate model to learn to predict y_t^{nj} by feeding it all available information from the other cities. However, this

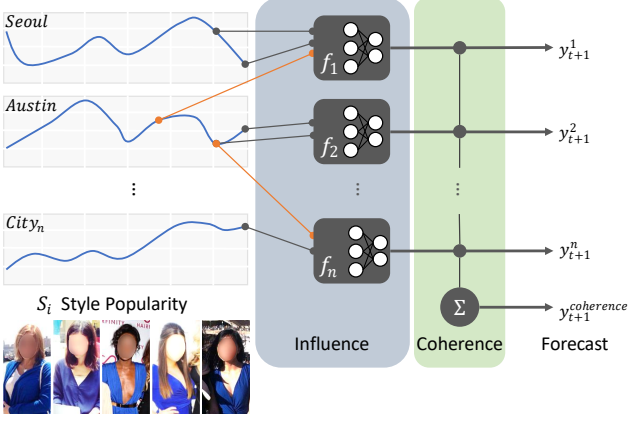


Figure 3: Influence coherent forecaster. Our model captures influence relations between cities for a given fashion style (orange connections) and uses them to predict future changes in the style popularity in each location. Additionally, our model regularizes the forecasts to be coherent with the global trend of the style observed across all cities.

approach does not satisfy the second property for an influencer since it does not constrain the influencer to have novel information that is not present in the influenced. Instead, our fashion influence relations can be captured using the Granger causality test [9]. The test determines that a time series y^1 Granger causes a time series y^2 if, while taking into account the past values of $y_{1..t}^2$, the past values of $y_{1..t}^1$ still have statistically significant impact on predicting the next value of y_{t+1}^2 . The test proceeds by modeling y^2 with an autoregressor of order d , *i.e.*:

$$y_t^2 = \phi_0 + \sum_{k=1}^d \phi_k y_{t-k}^2 + \sigma_t, \quad (2)$$

where σ_t is an error term and ϕ_k contains the regression coefficients. Then the autoregressor of y^2 is extended with lagged values of y^1 such that:

$$y_t^2 = \phi_0 + \sum_{k=1}^d \phi_k y_{t-k}^2 + \sum_{l=m}^q \psi_l y_{t-l}^1 + \sigma_t. \quad (3)$$

If the extended lags from y^1 do add significant explanatory power to y_t^2 , *i.e.* the forecast accuracy of y^2 is significantly better ($p < 0.05$) according to a regression metric (mean squared error), then y^1 Granger causes y^2 .

We estimate the influence relations across all cities' trajectories for each fashion style S^i . (In experiments we consider 50 such styles, and consider lags ranging from 1 to 8 temporal steps.) In this way, we establish the influence relations among cities and at which lag this influence occurs.

3.3. Coherent Style Forecaster

After we estimate the influence relations across the cities, we build a forecaster for each trajectory y^{ij} such that:

$$\hat{y}_{t+1}^{ij} = f(L(y_t^{ij}), I(y_t^{ij})|\theta), \quad (4)$$

where $I(y_t^{ij})$ is the set of lags from the influencer of y^{ij} relative to time step t as determined in the previous section, and $L(y_t^{ij})$ are the lags from y^{ij} 's own style popularity trajectory. We model $f(\cdot)$ using a multilayer perceptron (MLP) and estimate the parameters θ by minimizing the mean squared error loss:

$$\mathcal{L}_{forecast} = \sum_t (y_{t+1}^{ij} - f(L(y_t^{ij}), I(y_t^{ij})|\theta))^2, \quad (5)$$

where y_{t+1}^{ij} is the ground truth value of y^{ij} at time $t + 1$.

Coherence loss Our forecast model in its previous form does not impose any constraints on the forecasted values in relation to each other. However, while we are forecasting the style popularity at each individual location given the influence from the others, the forecasted popularities ($y_t^{i1}, y_t^{i2} \dots y_t^{in}$) are still for a common fashion style S^i that by itself exhibits a worldwide trend across all locations.

We propose to reconcile the base forecasts produced at each location through a *coherence loss* that captures the global trend. For all forecasts \hat{y}_{t+1}^{ij} for a fashion style S^i and across all cities $C^j \in C$, we constrain the distribution mean of the predicted values to match the distribution mean of the ground truth values:

$$\mathcal{L}_{coherence} = \frac{1}{|C|} \left(\sum_k y_{t+1}^{ik} - \sum_k \hat{y}_{t+1}^{ik} \right). \quad (6)$$

The coherence loss, in addition to capturing the global trend of S^i , helps in combating noise at the trajectory level of each city through regularizing the mean distribution of all forecasts. The final model is trained with the combined forecast and coherence losses:

$$\mathcal{L} = \mathcal{L}_{forecast} + \lambda \mathcal{L}_{coherence}. \quad (7)$$

Fig. 3 illustrates our model. For a style S_i , we model its popularity trajectory at each location (*e.g.*, Seoul) with a neural network of 2 layers and sigmoid non-linearity. The input of the network is defined by the lags from its own trajectory (shown in black) and any other influential lags from other cities discovered by the previous step (Sec. 3.2) which are shown in orange. Furthermore, the output of all local forecasters is further regularized to be coherent with the overall observed trend of S_i using our coherence loss.

4. Evaluation

In the following experiments, we demonstrate our model's ability to forecast styles' popularity changes

around the globe using the discovered influence relations. Furthermore, we analyze the influence patterns revealed by our model between major cities, how they influence global trends, and their dynamic influence trends through time.

Dataset We evaluate our approach on the GeoStyle dataset [28] which is based on Instagram and Flickr photos showing people from 44 major cities from around the world. In total, the dataset has 7.7 million images that span a time period from July 2013 until May 2016. The dataset is used for research purposes only.

Styles and popularity We use attribute predictions from [28] to represent each photo with 46 fashion attributes (e.g. colors, patterns and garment types). Based on these, we learn 50 fashion styles using a Gaussian mixture model. Then, for each style we infer its popularity trajectories in each city using a temporal resolution of weeks (cf. Sec. 3.1). While we focus on fashion styles in this work (each of which aggregates an array of commonly co-occurring attributes), we find similar results when considering individual visual attributes as the fashion concept of interest.

Implementation details We set $\lambda = 1$ for the coherence loss weight (see Eq. 7) and optimize our neural influence model using Adam for stochastic gradient descent with a learning rate of 10^{-2} and l_2 weight regularization of 10^{-8} . We select the best model based on the performance on a disjoint validation split using early stopping.

4.1. Influence-based Forecasting

We evaluate how well our model produces accurate forecasts by leveraging the influence pattern, and compare it to several baselines and existing methods [1, 28] that model trajectories in isolation. We adopt a *long-term forecasting* setup where we use the last 26 points from each style trajectory for testing, and the rest for model training. We arrange the baselines into three main groups:

Naive models: these models rely on basic statistical properties of the trajectory to produce a forecast. We consider five variants of these baselines similar to those from [1]; see supplementary for a detailed description.

Per-City models: These fit a separate parametric model trained on the history “lags” of each of the trajectories [4].

- FashionForward (EXP) [1]: an exponential decay model which forecasts based on a learned weighted average of the historical values.

- AR: it fits a standard autoregression model.

- ARIMA: the standard autoregressive integrated moving average model.

- GeoModel [28]: a parametric seasonal forecaster.

To our knowledge, the two existing methods [1, 28] represent the only prior vision approaches for style forecasting. Further, unlike our approach, all of the per-city models

Model	Seasonal		Deseasonalized	
	MAE	MAPE	MAE	MAPE
Naive				
Gaussian	0.1301	33.23	0.1222	26.08
Seasonal	0.0925	22.64	0.1500	33.39
Mean	0.0908	23.57	0.0847	18.97
Last	0.0893	22.20	0.1053	23.08
Drift	0.0956	23.65	0.1163	25.32
Per City				
FashionForward (EXP) [1]	0.0779	19.76	0.0848	18.94
AR	0.0846	21.88	0.0846	18.95
ARIMA	0.0919	23.70	0.1033	22.70
GeoModel [28]	0.0715	17.86	0.0916	20.31
All Cities				
VAR	0.0771	19.25	0.0929	20.41
Ours – Influence-based				
Full	0.0699	17.38	0.0824	18.29
w/o Influence	0.0708	17.70	0.0859	19.24
w/o Influence & Coherence	0.0858	20.95	0.0942	20.62

Table 1: Forecast errors on the GeoStyle dataset [28] for seasonal and deseasonalized fashion style trajectories.

consider the popularity trajectories of the styles in isolation, *i.e.*, they do not take into consideration possible interactions among the cities.

All-Cities models: These fit a parametric model trained on the trajectories of a style across all cities. Such models assume a full and simultaneous interaction between all cities. We consider the VAR model [4] to represent this group.

We compare all models using the forecast error captured by the mean absolute error (MAE), which measures the absolute difference between the forecasted and ground truth values, and the mean absolute percentage error (MAPE), which measures the forecast error scaled by the ground truth values. Additionally, to quantify the impact of possible seasonal yearly trends in fashion styles, we also consider forecasting the deseasonalized style trajectories (*i.e.* we subtract the yearly seasonal lag from the trajectories). The deseasonalized test is interesting because it requires methods to capture the subtle visual trends not simply associated with the location’s weather and annual events.

Table 1 shows the results. The proposed model outperforms all the naive, per-city, and all-cities models. This shows the value of discovering influence for the quantitative forecasting task. Ablation studies (bottom segment of the table) show the impact of each component of our model. We evaluate two versions of our model, one without our influence modeling from Sec. 3.2, which assumes a full interaction pattern among all cities, and a second version that also is not trained for coherent forecasts (Sec. 3.3). We see that both our influence modeling and coherence-based forecasts are important for accurate predictions.

We notice that the styles’ popularity trajectories do have a seasonal component; seasonal models (like GeoModel [28] and Seasonal) do well compared to non-

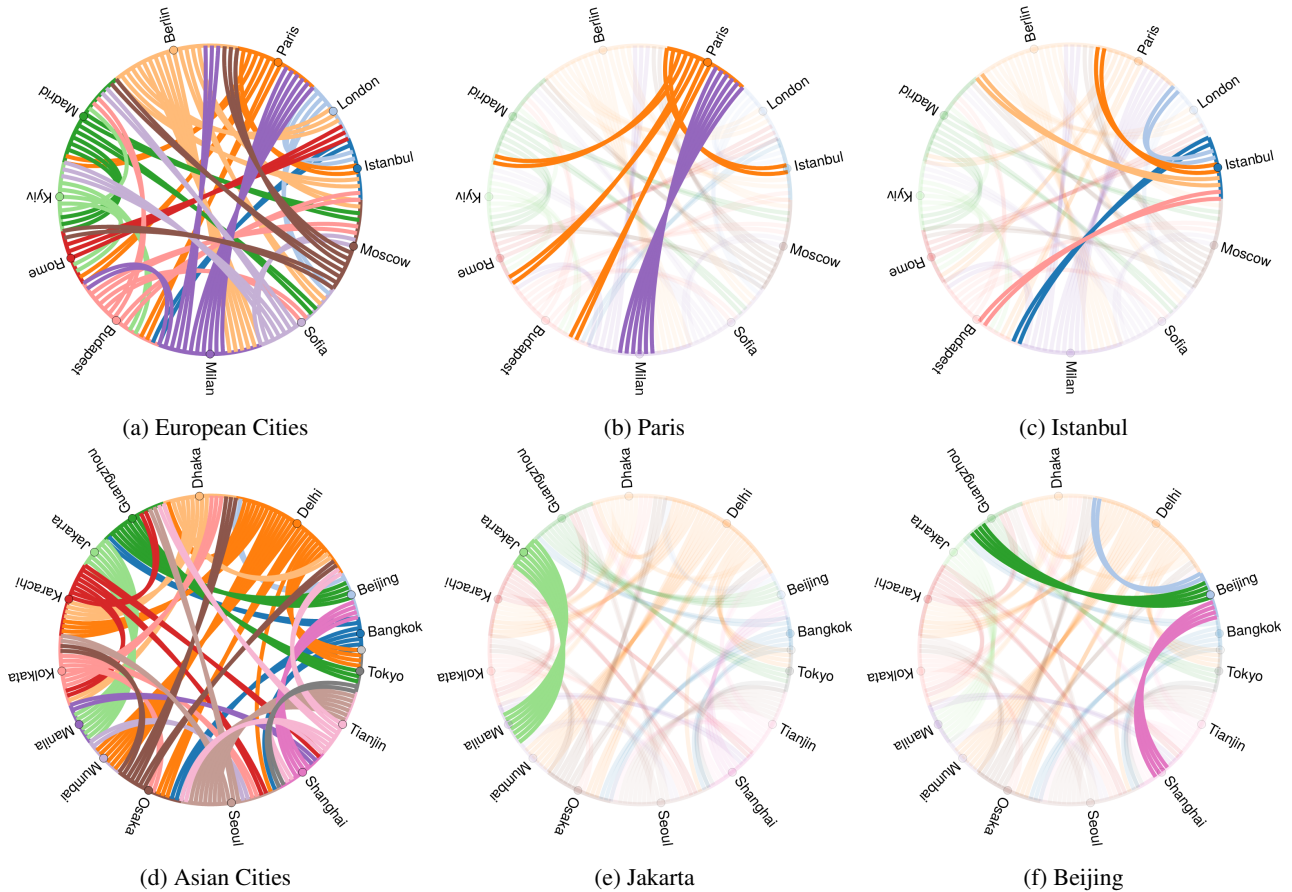


Figure 4: Style influence relations discovered by our model among European (a) and Asian (d) cities. The number of chords coming out of a node (*i.e.* a city) is relative to the influence weight of that city on the receiver. Chords are colored according to the source node color, *i.e.* the influencer. Our model discovers various types of influence relations from multi-city (*e.g.* Paris) and single-city (*e.g.* Jakarta) influencers to cities that are mainly influence receivers (*e.g.* Istanbul and Beijing).

seasonal ones (like AR and Drift), but still underperform our approach. This ranking changes on the deseasonalized test, where models like FashionForward [1] and AR outperform the seasonal ones. Our model outperforms all competitors on both types of trajectories, which demonstrates the benefits of accounting for influence. Our model’s improvements compared to the best “non-influence” per-city or all-cities competitors on both types of style trajectories are statistically significant based on a t-test with $p < 0.05$.

4.2. Influence Relations

The results thus far confirm that our method’s discovered influence patterns are meaningful, as seen by their positive quantitative impact on forecasting accuracy. Next, we analyze them qualitatively to understand more about what was learned. We consider influence interactions along two axes: 1) a local one that looks at pairwise influence relations among the cities; and 2) a global one which examines how cities influence the world’s fashion trends.

1) City \rightarrow City influence For each visual style, our model estimates the influence relation between cities and at which temporal lag, yielding a tensor $\mathbf{B} \in \mathbb{R}^{|C| \times |C| \times |S|}$ such that B_{ij}^k is the influence lag of C^i on C^j for style S^k . By averaging these relations across all visual styles, we get an estimate of the overall influence relation between all cities weighted by the temporal length, *i.e.* long term influencers are given more weight than instantaneous ones. We visualize the influence relations using a directed graph where each node represents a city, and we create a weighted edge from city C^i to C^j if C^i is found to be influencing C^j .

Fig. 4 shows an example of the influence pattern for fashion styles discovered by our model among major European (top left) and Asian (bottom left) cities, where the number of connections between two cities is relative to the weight of the influence relation. Our model discovers interesting patterns. For example, there are a few fashion hubs like *Paris* and *Berlin* which exert influence on multiple cities while at the same time being influenced by few (one or two)

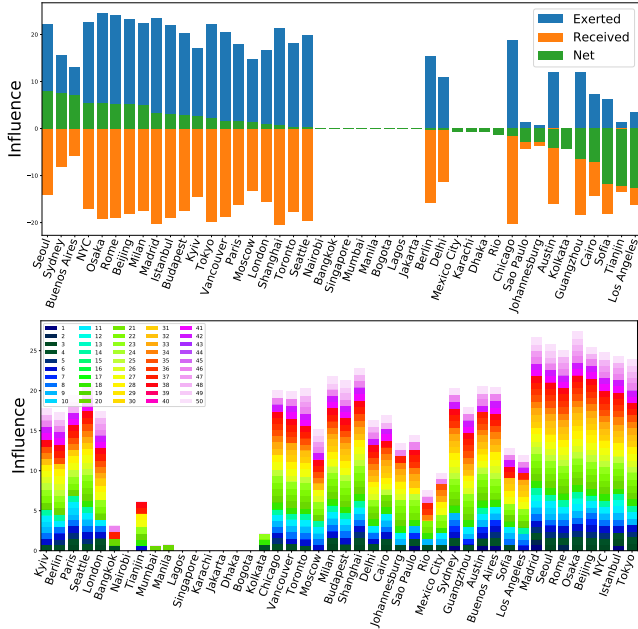


Figure 5: Top: Worldwide ranking of major cities according to their fashion influence on their peers. The cities are sorted by their *net* influence score (green). The center cities with no bars indicate that they do not have influential relations with above average weight. Bottom: the exerted influence score is split into influence score per individual style for each city (sorted by style influence similarity).

cities. *Paris* influences 4 cities in Europe while being influenced by *Milan* only (Fig. 4b). Cities like *Jakarta* have a one-to-one influence relation with *Manila* (Fig. 4e). On the other end of the spectrum, we find cities like *Istanbul* (Fig. 4c) and *Beijing* (Fig. 4f) that mainly receive influence from multiple sources while influencing few.

Additionally, we rank all cities according to their accumulated influence power on their peers. That is, we assign an influence score for each city according to the sum of weighted influence relations *exerted* by that city on the rest. Similarly, we also calculate the sum of *received* influence as well as the difference in both as the *net* influence score. Fig. 5 (top) shows these three influence scores for all cities across the world, sorted by the *net* score. The ranking reveals that some cities, like *London* and *Seattle*, act like focal points for fashion styles, *i.e.* they receive and exert a high volume of influence simultaneously. Others, like *Seoul* and *Osaka*, have a high net influence, which could indicate having some unique fashion styles not influenced by external players. Furthermore, breaking down the exerted influence score for each city to per-style influence scores, we see in Fig. 5 (bottom) that we can identify and group influencers into “teams” based on their common set of styles where they exert their influence. For example, *Chicago*, *Vancouver*,

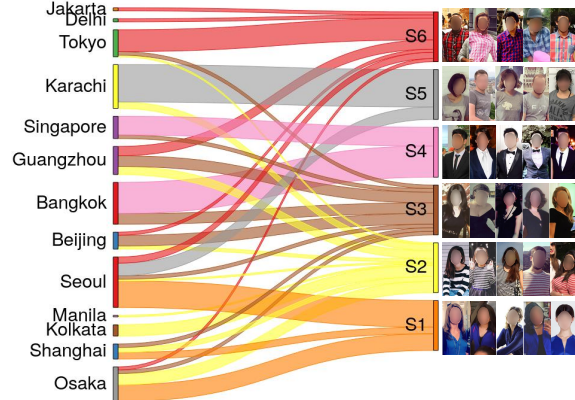


Figure 6: Discovered influence by our model of Asian cities on global trends of 6 fashion styles. The width of the connection is relative to the influence weight of that city in relation to other influencers of the same style.

ver, and *Toronto* constitute a team since they seem to be influencing similar sets of visual styles.

2) City → World influence Alternatively, we can analyze the influence relation between a city and the world trend for a specific style. This helps us better understand who are the main influencers on the world stage for each of the styles. We capture this relation by modeling the interaction of a city’s popularity trajectory on the global one (*i.e.* the observed trend of the collective popularity of the same style across the world).

Fig. 6 shows a set of Asian cities and their influence on six global fashion styles. We see that for some of the fashion styles, like S_4 and S_5 , a couple of cities maintain a monopoly of influence on them, whereas others, like S_3 and S_2 , are influenced almost uniformly by multiple cities. Our influence model also reveals the influence strength (measured by the temporal lag) of these cities relative to their peers at the world stage. See for example the strong influence of *Seoul* and *Bangkok* compared to the delicate one of *Manila* and *Jakarta*, as represented by the width displayed for their respective influence relation to the global trends.

4.3. Influence Correlations

Next we analyze the correlations of these relations discovered by our model with known real-world properties of the cities. We stress that the trends visible in the photos *are* exactly what our model measures; there is no separate “ground truth” against which to score the influence measurements. Correlating against other properties simply helps unpack what the trends do or do not relate to. We collect information about the annual gross domestic product (GDP), the geolocation, the population size, and the yearly average temperature for each of the cities. We calculate the correlations of these properties with the influence informa-

Meta Info.	Fashion Influence	
	Direction	World Rank
GDP	0.037	0.373
Temperature	- 0.319	- 0.616
Latitude	- 0.348	0.596
Population	0.038	- 0.193
Distance	- 0.165	n/a
Num. Samples	- 0.148	0.086

Table 2: Correlations of the discovered influence patterns with meta information about the cities.

tion discovered by our model at two levels: 1) influence world ranking (*i.e.* does a high influence rank correlate with the population size of the city? do cities with warm weather have a high influence score?), and 2) relation direction (*i.e.* does influence flow from high to low GDP cities? do cities influence those that are geographically close to them?).

World rank In Table 2 (2nd column) we correlate the discovered influence ranking of all the cities with the ranking derived from each of the meta properties using the Spearman coefficient. The correlation of these meta properties with the overall ranking of the cities uncovers some curious cases. There is an above average correlation between the city influence rank and its latitude; many of the influential fashion cities are on the northern hemisphere. We see a weaker but positive correlation with GDP, *i.e.* a higher GDP could be a faint indicator of a higher fashion influence. Finally, we observe a negative correlation with average temperature; influential fashion cities are often colder.

Relation direction Table 2 (1st column) shows the correlation of the influence directions discovered by our model with differences in each of the meta properties between the influencer and the influenced city. Specifically, for each city and meta property M_i (*e.g.* GDP), first we measure the differences between that city and the rest in regards to M_i , then we correlate these differences with the influence exerted by that city. Interestingly, none show high correlation with fashion influence directions. The relation type cannot be reliably estimated based on the differences in GDP (*e.g.* high GDP cities do not always influence lower GDP ones), population (*e.g.* cities with high population do not necessarily influence others with lower population or vice versa), nor distance (*e.g.* influence does not correlate well with how far one city is from its influencer). A weak and negative correlation is found with temperature and latitude differences, showing that cities with similar temperature or at similar latitudes tend to influence each other slightly more. These results suggest our model discovers complex fashion influence relations that are hard to infer from generic properties of the constituent players.

As a sanity check, we also explore the correlation of the




Figure 7: Dynamics analysis of exerted fashion influence at multiple time steps (with a 3 month interval) reveals the cities’ temporal changes in influence strength.

number of image samples collected from each city in the dataset with the two types of influence information (Table 2 last row). We find that there is no strong correlation between the learned influences and the number of images available in the data for each city (*i.e.* influential cities are not those with a higher number of samples in the dataset).

4.4. Influence Dynamics

Finally, we study the changes in the influence rank of the cities through time. We carry out our influence modeling based on the style trajectories of the various cities as before, but at multiple sequential time steps. Then we collect the overall influence score of each city at each step.

Fig. 7 shows the change in the influence score for a subset of 5 cities spanning different continents. We notice that cities show various dynamic behaviors across time. While some cities like *London* and *Rio* maintain a steady influence score through time (at different levels), others like *Austin* and *Johannesburg* demonstrate a positive trend and are gaining more influence in the fashion domain over time but at varying speeds. Other cities like *Tianjin* exhibit a mild decline in their fashion influence.

5. Conclusion

We introduced a model to quantify influence of visual fashion trends, capturing the spatio-temporal propagation of styles around the world. Our approach integrates both influence relations and a coherence regularizer to predict future style popularity conditioned on place. Both our forecasting results and our analysis of the learned influences suggest potential applications in social science, where computer vision can unlock trends that are otherwise hard to capture.

Acknowledgements: We thank Utkarsh Mall for helpful input on the GeoStyle data. UT Austin is supported in part by NSF IIS-1514118.

References

- [1] Ziad Al-Halah, Rainer Stiefelhagen, and Kristen Grauman. Fashion Forward: Forecasting Visual Style in Fashion. In *ICCV*, 2017.
- [2] Tamara L Berg, Alexander C Berg, and Jonathan Shih. Automatic attribute discovery and characterization from noisy web data. In *ECCV*, 2010.
- [3] Lukas Bossard, Matthias Dantone, Christian Leistner, Christian Wengert, Till Quack, and Luc Van Gool. Apparel classification with style. In *ACCV*, 2012.
- [4] George EP Box, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.
- [5] Chengyao Chen, Zhitao Wang, Wenjie Li, and Xu Sun. Modeling scientific influence for research trending topic prediction. In *AAAI*, 2018.
- [6] Huizhong Chen, Andrew Gallagher, and Bernd Girod. Describing clothing by semantic attributes. In *ECCV*, 2012.
- [7] Kuanting Chen, Kezhen Chen, Peizhong Cong, Winston H Hsu, and Jiebo Luo. Who are the Devils Wearing Prada in New York City? In *ACM Multimedia*, 2015.
- [8] Haoye Dong, Xiaodan Liang, Xiaohui Shen, Bochao Wang, Hanjiang Lai, Jia Zhu, Zhiting Hu, and Jian Yin. Towards multi-pose guided virtual try-on network. In *ICCV*, 2019.
- [9] Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, pages 424–438, 1969.
- [10] M Hadi Kiapour, Xufeng Han, Svetlana Lazebnik, Alexander C Berg, and Tamara L Berg. Where to buy it: Matching street clothing photos in online shops. In *ICCV*, 2015.
- [11] M Hadi Kiapour and Robinson Piramuthu. Brand > logo: Visual analysis of fashion brands. In *ECCV Workshops*, 2018.
- [12] Xintong Han, Zuxuan Wu, Weilin Huang, Matthew R. Scott, and Larry S. Davis. Finet: Compatible and diverse fashion image inpainting. In *ICCV*, 2019.
- [13] Ruining He and Julian McAuley. Ups and Downs: Modeling the Visual Evolution of Fashion Trends with One-Class Collaborative Filtering. In *WWW*, 2016.
- [14] Shintami Chusnul Hidayati, Cheng-Chun Hsu, Yu-Ting Chang, Kai-Lung Hua, Jianlong Fu, and Wen-Huang Cheng. What dress fits me best?: Fashion recommendation on the clothing style for personal body shape. In *ACM Multimedia*, 2018.
- [15] Wei-Lin Hsiao and Kristen Grauman. Learning the latent look: Unsupervised discovery of a style-coherent embedding from fashion images. In *ICCV*, 2017.
- [16] Wei-Lin Hsiao and Kristen Grauman. Creating capsule wardrobes from fashion images. In *CVPR*, 2018.
- [17] Wei-Lin Hsiao and Kristen Grauman. Dressing for diverse body shapes. In *CVPR*, 2020.
- [18] Wei-Lin Hsiao, Isay Katsman, Chao-Yuan Wu, Devi Parikh, and Kristen Grauman. Fashion++: Minimal edits for outfit improvement. *ICCV*, 2019.
- [19] Junshi Huang, Rogerio S Feris, Qiang Chen, and Shuicheng Yan. Cross-domain image retrieval with a dual attribute-aware ranking network. In *ICCV*, 2015.
- [20] Hirokatsu Kataoka, Yutaka Satoh, Kaori Abe, Munetaka Minoguchi, and Akio Nakamura. Ten-million-order human database for world-wide fashion culture analysis. In *CVPR Workshops*, 2019.
- [21] M Hadi Kiapour, Kota Yamaguchi, Alexander C Berg, and Tamara L Berg. Hipster wars: Discovering elements of fashion styles. In *ECCV*, 2014.
- [22] Gunhee Kim and Eric Xing. Discovering pictorial brand associations from large-scale online image data. In *ICCV Workshops*, 2013.
- [23] Zhanghui Kuang, Yiming Gao, Guanbin Li, Ping Luo, Yimin Chen, Liang Lin, and Wayne Zhang. Fashion retrieval via graph reasoning networks on a similarity pyramid. In *ICCV*, 2019.
- [24] Iljung S Kwak, Ana Cristina Murillo, Peter N Belhumeur, David J Kriegman, and Serge J Belongie. From bikers to surfers: Visual recognition of urban tribes. In *BMVC*, 2013.
- [25] Si Liu, Jiashi Feng, Zheng Song, Tianzhu Zhang, Hanqing Lu, Changsheng Xu, and Shuicheng Yan. Hi, magic closet, tell me what to wear! In *ACM Multimedia*, 2012.
- [26] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *CVPR*, 2016.
- [27] Zheng Lu and Kristen Grauman. Story-driven summarization for egocentric video. In *CVPR*, 2013.
- [28] Utkarsh Mall, Kevin Matzen, Bharath Hariharan, Noah Snavely, and Kavita Bala. Geostyle: Discovering fashion trends and events. In *ICCV*, 2019.
- [29] Kevin Matzen, Kavita Bala, and Noah Snavely. Streetstyle: Exploring world-wide clothing styles from millions of photos. *arXiv preprint arXiv:1706.01869*, 2017.
- [30] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. Image-based recommendations on styles and substitutes. In *ACM SIGIR Conference on Research and Development in Information Retrieval*, 2015.
- [31] Rishabh Misra, Mengting Wan, and Julian McAuley. Decomposing fit semantics for product size recommendation in metric spaces. In *ACM Conference on Recommender Systems*, 2018.
- [32] Ana C Murillo, Iljung S Kwak, Lubomir Bourdev, David Kriegman, and Serge Belongie. Urban tribes: Analyzing group photos from a social perspective. In *CVPR Workshops*, 2012.
- [33] Dafna Shahaf and Carlos Guestrin. Connecting the dots between news articles. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2010.
- [34] Edgar Simo-Serra, Sanja Fidler, Francesc Moreno-Noguer, and Raquel Urtasun. Neuroaesthetics in Fashion: Modeling the Perception of Fashionability. In *CVPR*, 2015.
- [35] Zheng Song, Meng Wang, Xian-sheng Hua, and Shuicheng Yan. Predicting occupation via human clothing and contexts. In *ICCV*, 2011.
- [36] Andreas Veit, Balazs Kovacs, Sean Bell, Julian McAuley, Kavita Bala, and Serge Belongie. Learning visual clothing style with heterogeneous dyadic co-occurrences. In *ICCV*, 2015.

- [37] Sirion Vittayakorn, Kota Yamaguchi, Alexander C. Berg, and Tamara L. Berg. Runway to realway: Visual analysis of fashion. In *WACV*, 2015.
- [38] Bochao Wang, Huabin Zheng, Xiaodan Liang, Yimin Chen, Liang Lin, and Meng Yang. Toward characteristic-preserving image-based virtual try-on network. In *ECCV*, 2018.
- [39] Wenguan Wang, Yuanlu Xu, Jianbing Shen, and Song-Chun Zhu. Attentive fashion grammar network for fashion landmark detection and clothing category classification. In *CVPR*, 2018.
- [40] Cong Yu, Yang Hu, Yan Chen, and Bing Zeng. Personalized fashion design. In *ICCV*, 2019.
- [41] Weijiang Yu, Xiaodan Liang, Ke Gong, Chenhan Jiang, Nong Xiao, and Liang Lin. Layout-graph reasoning for fashion landmark detection. In *CVPR*, 2019.
- [42] Bo Zhao, Jiashi Feng, Xiao Wu, and Shuicheng Yan. Memory-augmented attribute manipulation networks for interactive fashion search. In *CVPR*, 2017.