

Temporal Analysis of Social Media Response to live events: the Milano Fashion Week

Marco Brambilla, Stefano Ceri, Florian Daniel, Gianmarco Donetti

Politecnico di Milano,
Piazza Leonardo da Vinci 32,I-20133 Milan, Italy
name.surname@polimi.it

Abstract. Social media response to catastrophic events, such as natural disasters or terrorist attacks, has received a lot of attention. However, social media are also extremely important in the context of planned events, such as fairs, exhibits, festivals, as they play an essential role in communicating them to fans, interest groups, and the general population. These kinds of events are geo-localized within a city or territory and are scheduled within a public calendar. We consider a specific scenario, the Milano Fashion Week (MFW), which is an important event in our city. We focus our attention on the spreading of social content in *time*, measuring the delay of the event propagation. We build different clusters of stakeholders (fashion brands), we characterize several features of time propagation and we correlate it to the *popularity* of involved actors. We show that the clusters by time and popularity are loosely correlated, and therefore the time response cannot be easily inferred. This motivates the development of a predictor through supervised learning in order to anticipate the space cluster of a new brand.

1 Introduction

Thanks to the wide adoption of smartphones, which enable continuous sharing of information with our social network connections, the online response to popular real world events is becoming increasingly significant, not only in terms of volumes of contents shared in the social network itself, but also in terms of velocity in the spreading of the news about events with respect to the time and to the geographical dimension. It has been noted that social signals are at times faster than media news with highly impacting events, such as terrorist attacks or natural disasters.

This work deals more specifically with the problem of social media response to a scheduled and popular real world event, the Milano Fashion Week occurred from the 24th to the 29th February of 2016, analysing the behaviour of users who re-acted (or pro-acted) in relationship with each specific fashion show during the week.

MFW, established in 1958, is part of the global “Big Four fashion weeks”, the others being located in Paris, London and New York [4]. The event is organized by *Camera Nazionale della Moda Italiana*, who manages and fully co-ordinates

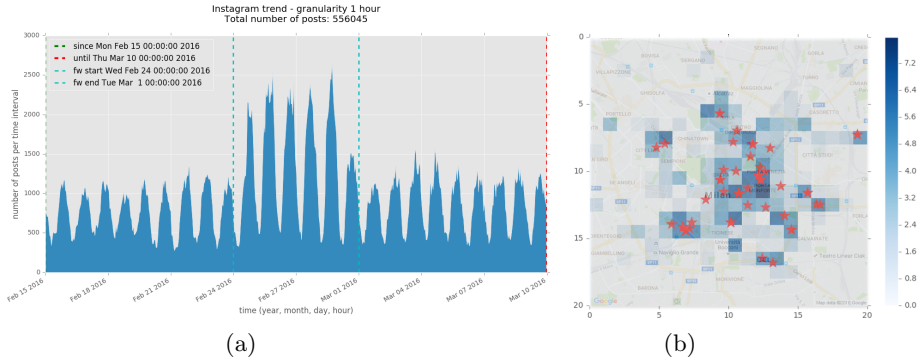


Fig. 1. Temporal overview for the three analyzed weeks of Instagram posts (a) and map representing the geographical distribution of MFW events (represented by red stars) and post density (b) - showing the spreading in the city of both the MFW events and the accumulated Instagram posts.

about 170 shows, presentations and events, thus facilitating the work of show-rooms, buying-offices, press offices, and public relations firms. Camera Nazionale della Moda carries out essential functions like drawing up the calendar of the shows and presentations, managing the relations with the institutions, the press office and creation of special events. MFW represents the most important meeting between market operators in the fashion industry. Out of the 170 shows, we are interested only in the catwalk shows, which are the core of the fashion week. The whole set of catwalks includes a total of 73 brands; among them, 68 brands organize one single event, 4 brands organize 2 events, and 1 brand organizes 3 events.

We formulate our problem as the analysis of the response in time of scheduled, popular events on social media platforms. Our goal is to describe and characterize the time at which social media respond to the events which appear in the official calendar and are linked to specific brands. Informally, we observe either peaks of reactions which then quickly disappear, or instead slower reactions that however tend to remain observable for a longer time. Estimating the time latency of social responses to events is important for the brands, which could plan reinforcement actions more accurately, essentially by adding well-planned social actions so as to sustain their social presence over time.

Our approach is as follows. We start by establishing correspondences between social events and calendarized events by filtering them with the same brand identity. We then discuss the Granger causality between the two time series. We next normalize and discretize the social media response curves, and finally we build the clustering of the brands.

We then cluster brands by online popularity, and show that popularity does not precisely explain the difference in time reactions, although it is weakly associated with it. This motivates the development of a predictor of the time dispersion of the social signal of a brand which is based just on time-specific features; our

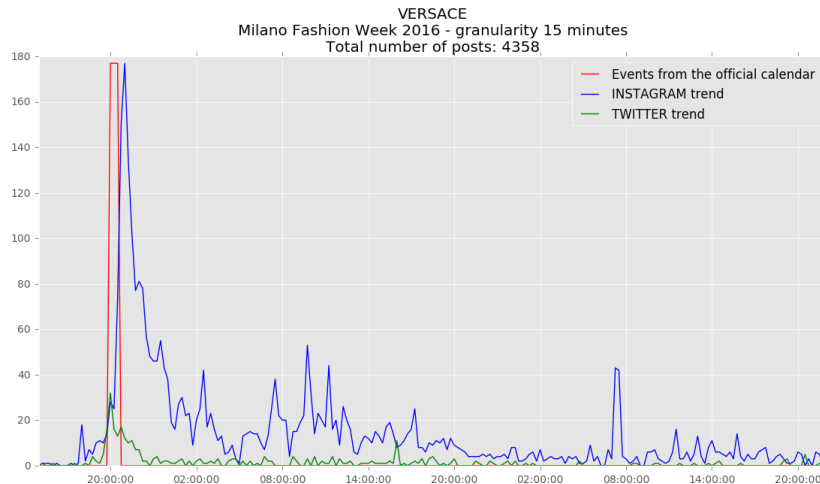


Fig. 2. Social media response to Versace’s event of 26th February at 20:00. The granularity is 15 minutes. The red line represents the *calendarSignal* function and the other lines report the number of posts in the specific time window. The blue line is for Instagram, the green one is for Twitter.

predictor uses supervised learning in order to anticipate the cluster of a new, unclustered brand.

The paper is organized as follows: Sections 2 present our approach to data collection and preparation; then Section 3 describes the data analysis leading to clustering of brands by time responses. Section 4 introduces brand popularity and shows the best matching between brands clustered by time response and popularity; then, Section 5 introduces the time class predictor for a new brand, Section 6 presents related work, and Section 7 concludes.

2 Data Collection and Preparation

We initially extracted posts by invoking the social network APIs of Twitter and Instagram; for identifying the social reactions to MFW, we used a set of 21 hashtags and keywords provided by domain experts in the fashion sector, i.e., researchers of the *Fashion in Process* group (FIP) of Politecnico di Milano¹. We focused on 3 weeks, before, during and after the event. In this way, we collected 106K tweets (out of which only 6.5% geolocated) and 556K Instagram posts (out of which 28% geolocated); eventually, we opted for considering only Instagram posts, as they represent a much richer source for the particular domain of Fashion with respect to Twitter. Figure 1 shows the temporal and geographical distribution of the posts.

¹ <http://www.fashioninprocess.com/>

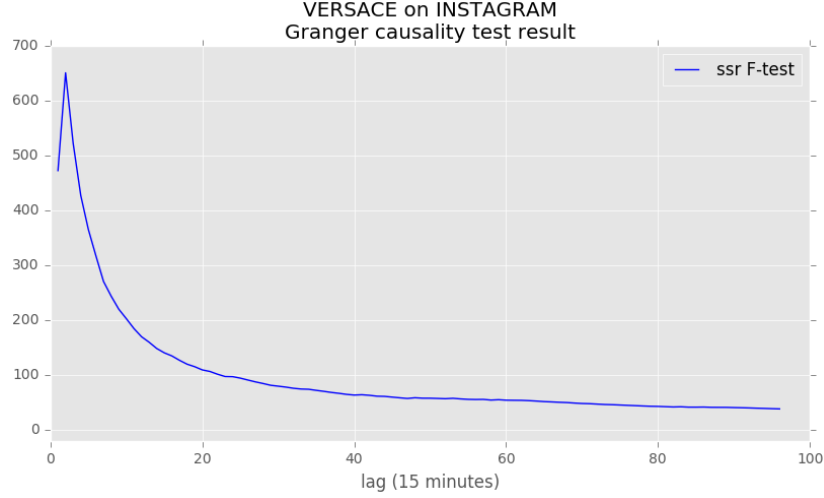


Fig. 3. Granger Causality test results for Versace between the *calendarSignal* and the number of Instagram posts in time.

We performed an initial analysis of the content, for associating each post with the corresponding event. In this specific scenario, the task was simple because each event was directly associated with a fashion brand, mentioned in the posts; the characterization of the brands was again provided by the FIP experts. For instance, for identifying the posts related to the Gucci catwalk, which was held on February 24th at 2:30pm in Milano, Via Valtellina 17, we collected the posts containing the hashtags and keywords *#Gucci* and *Gucci*, filtering the posts through suitable regular expressions. This allowed us to collect 7718 Instagram posts related to that specific event.

3 Time Response Analysis

For studying the temporal dynamics of the social media reactions, we consider two signals: the temporal series of number of posts related to a given brand B_i , and the one recording the presence of a live event for the brand B_i . The latter, which we name *calendar signal* for given brand B_i , is valid through the time intervals Δt of analysis and is defined as:

$$calendarSignal(B_i, \Delta t) = \sum_{e \in B_i} 1_e(\Delta t) \quad (1)$$

where $1_e(\Delta t)$ indicates the presence of the event e in the time window Δt . Intuitively, the signal has value 1 in the intervals when a live event of brand B_i is taking place, while it has value 0 when no events are taking place.

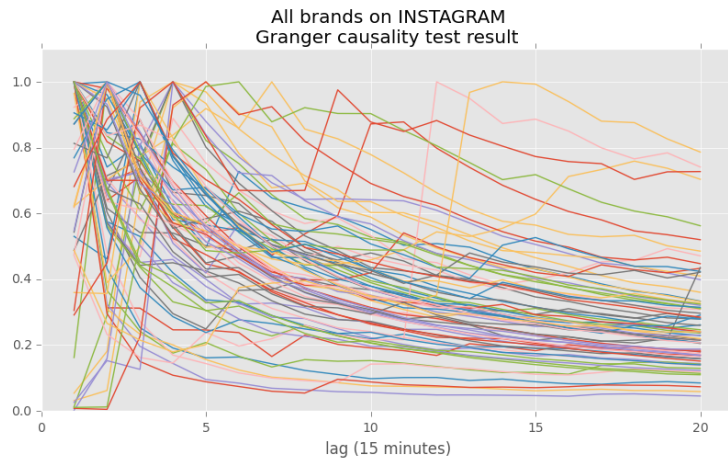


Fig. 4. Normalized Granger Causality test results for all the 65 analyzed brands.

3.1 Granger Causality

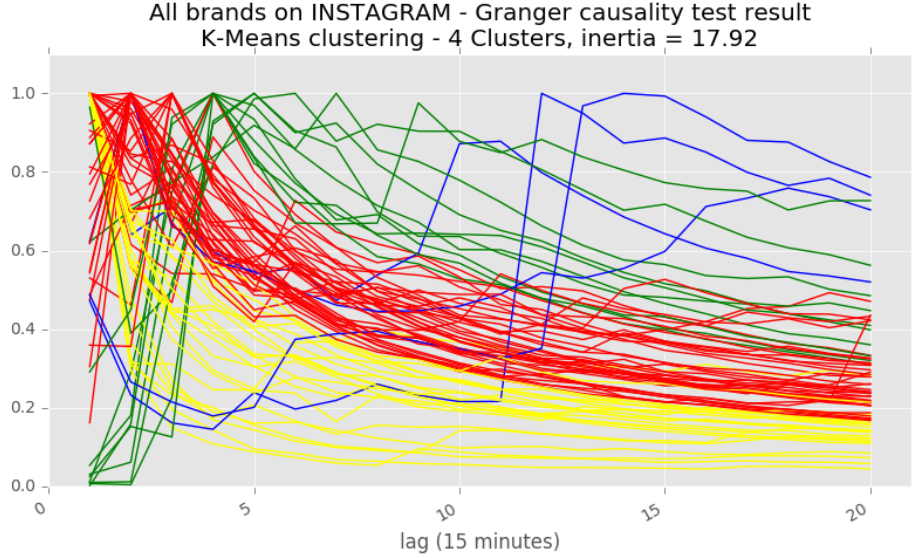
We then focused on determining a causality relationship between the events and the follow-up posts. Specifically, we measured the Granger causality between two different time series, (a) specific events and (b) posts reacting to those events; this had to be done brand by brand, by suitable selections both from the calendar and from the dataset of social posts.

The Granger Causality test applies to two different time series; a time series is said to Granger-cause another time series if it can be shown, usually through a series of t-tests and F-tests, that the values of the first series provide statistically significant information about future values of the second series. In particular, we want to reject, with statistical significance, the null hypothesis of the test that the calendar signal does not Granger-cause the social media signal. We focus on the F-test, which evaluates the ratio of two scaled sums of squares reflecting different sources of variability, constructed so that the statistic tends to be greater when the null hypothesis is not true.

As an example, Figure 2 shows the social response to the event of *Versace* on 26th February at 8pm. Note a strong reaction in the social media relatively close to the scheduled events: indeed, we have a peak of about 180 Instagram posts in the time window starting when the event is just completed, and then the number of posts per time window decreases rapidly. Figure 3 shows the graph of the Granger causality for Versace. We performed tests like these for all the brands that have one or more events in the Milano Fashion Week 2016 calendar, confirming Granger causality for all the brands. Figure 4 shows the normalized Granger Causality results for all the 68 analyzed brands.

Table 1. Information on the intra-cluster inertia and inertia gain when adding one additional cluster, after clustering the results with different values of k , for $3 \leq k \leq 5$.

k	Inertia	Inertia Gain
3	23.82	12.39
4	17.92	5.90
5	14.87	3.05

**Fig. 5.** Clusters produced by k-means clustering in 20-dimensional space with $k=4$.

3.2 Clustering

We then looked for similarity of the time response curves for posts. We normalized each peak at 1, since we were no longer interested in the statistical relevance, but just in the shape of the curves. We then considered a period of 5 hours and applied k-means clustering in L -dimensional space, where $L = 20$ is the number of points we have collected for each curve (therefore, we discretized the curve using 15 minutes intervals).

For deciding the ideal number k of clusters which better describes the scenario, we computed clustering for $k \leq 15$, and then we computed inertia for each such choice. Inertia (or within-cluster sum-of-squares) is a measure of internal coherence; the inertia curve is monotonically decreasing, with the maximum value corresponding to just one global cluster and the minimum value equal to 0 when the number of clusters coincide with the number of elements.

Based on inertia values (partially reported in Table 1), we picked k equal to 4. The values of inertia justify our choice, as: $k = 3$ corresponds to 23.82, $k = 4$

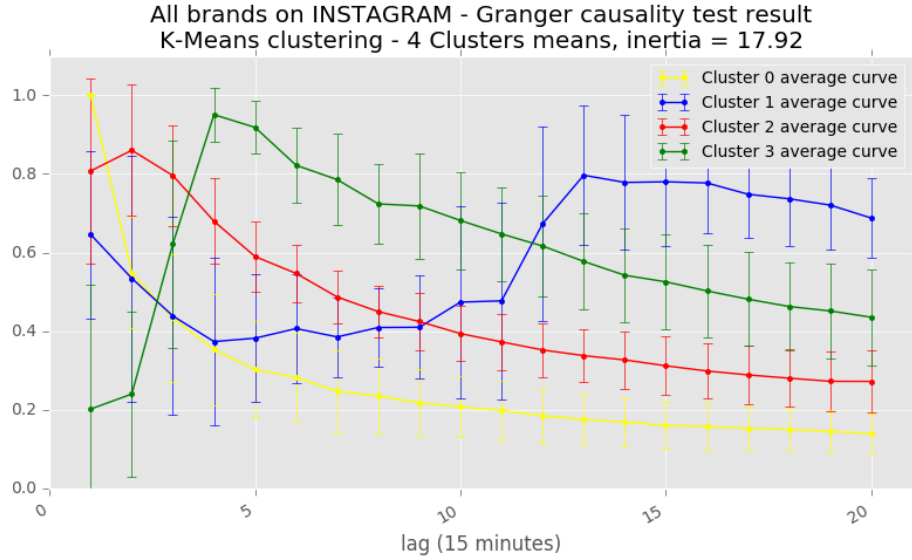


Fig. 6. Most representative elements of each cluster.

to 17.92, $k = 5$ to 14.87, and the inertia is decreasing very slowly for $k \geq 4$. The resulting clusters are associated to a color code:

- *Yellow*, with high immediate response;
- *Red*, with lagged response peak at 15 minutes;
- *Green*, with lagged response peak at 45-60 minutes;
- *Blue*, with an initial significant response but with another lagged response peak after 3 hours.

Figure 5 shows the clusters produced. We then selected the most representative elements for each cluster as the closest one to the centroid of each cluster; they are shown in Fig. 6:

- *Costume National*, from the yellow group;
- *Trussardi*, from the red group;
- *Alberta Ferretti*, from the green group;
- *Emporio Armani*, from the blue group.

4 Comparison with Brand Popularity

4.1 Popularity Analysis

We next turned to a simpler observation, the *brand popularity*, in order to evaluate the popularity of a brand and see if it relates to the above clusters; we focused on 65 brands which were hosting fashion shows during MFW. In our

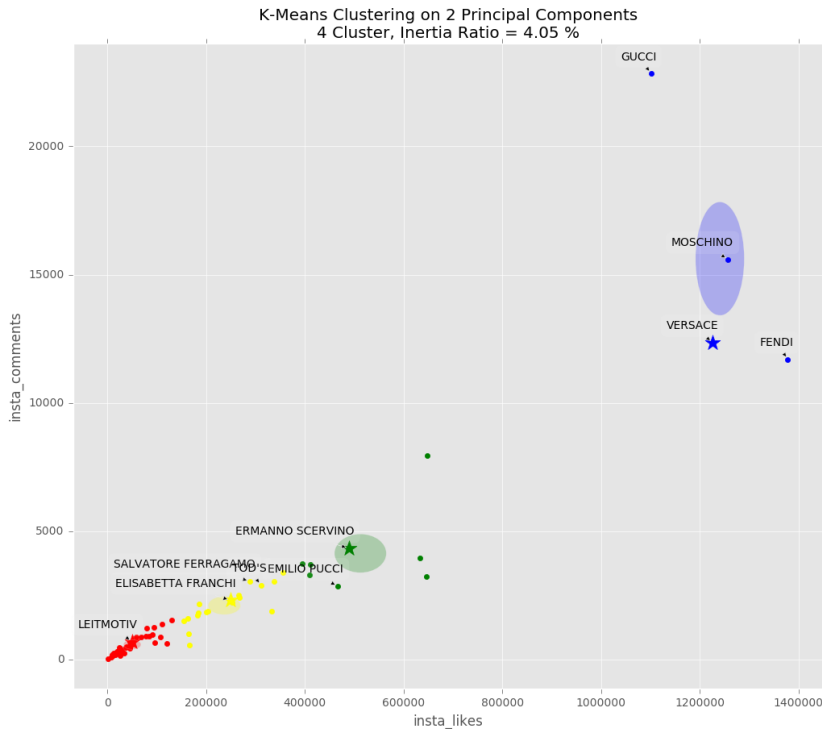


Fig. 7. Representation of popularity clusters with respect to the two principal components (comments and likes on Instagram, respectively).

analysis, We extracted from our Twitter and Instagram datasets a classic set of popularity features, related to each brand: the number of posts on Instagram, number of likes collected on Instagram, number of comments collected on Instagram, number of posts on Twitter, number of likes collected on Twitter, number of retweets collected on Twitter. We then performed a principal component analysis (PCA) in order to find the best features, and not surprisingly we noticed that likes on Instagram essentially dominate, as the 2 principal components are:

- Number of likes on Instagram (99.9% of total variance)
- Number of comments on Instagram (0.0025% of total variance)

In the end, we run k-means over these attributes, asking again for 4 different clusters, in order to better compare our final results. Figure 7 shows the outcome of this clustering. The groups could be described as following: from the red cluster to the blue cluster we are going from the most unpopular brands to the most popular ones, in the two social media of Twitter and Instagram. These results were confirmed by our experts in fashion design.

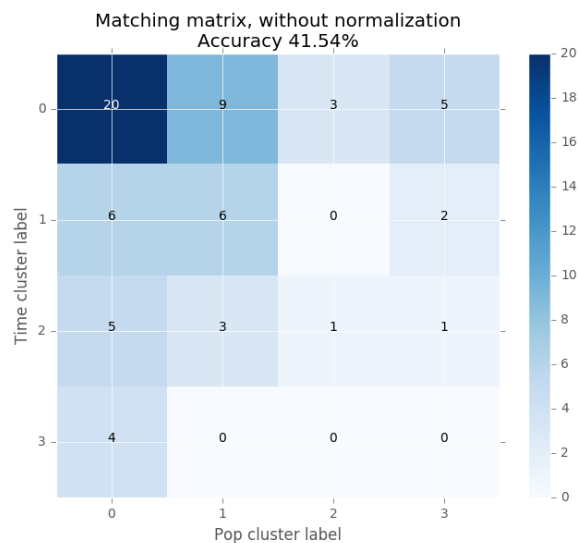


Fig. 8. Matching matrix comparing time response versus popularity response.

4.2 Time – Popularity Correlation

We then studied the correlation among the two clusterings. We label one clustering result arbitrarily and then rename the other clustering labels with all the possible permutations in the set of adopted labels. For each re-labelling, we compute a measure of correlation between the two clustering results, assuming a correspondence between same-name-labels, and we take the best renaming permutation in terms of the specific measure adopted. We pick as statistic measure of validation the accuracy in juxtaposing one cluster to another one. We recall that we are in a multiclass case, then the accuracy will be the percent of true-matching between the two clustering results.

We visualize the best matching by using a confusion matrix, a table that visualizes the correlations of the two different results clustering. One clustering is taken on the rows, the other one on the columns, total correlation requires all diagonal elements to be one and all other elements to be zero; with partial correlation we measure the accuracy of the matrix as the sum of all correlation values appearing along the diagonal. The best accuracy we obtain is 41.54%, that produces the confusion matrix in Figure 8. In a few words, the best correlation is obtained juxtaposing:

- The red cluster from time, with lagged response peak at 15 minutes, with the red cluster from popularity, the most unpopular ones;
- The yellow cluster from time, with high immediate response, with the yellow cluster from popularity, the third ones for popularity;
- The green cluster from time, with lagged response peak at 45-60 minutes, with the blue cluster from popularity, the most popular ones;

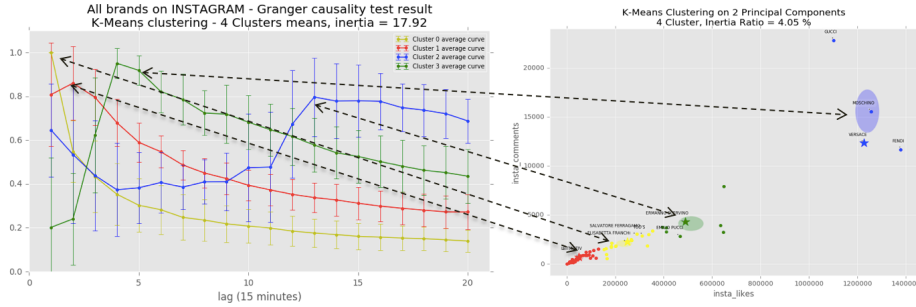


Fig. 9. Visual mapping between the temporal clusters and the popularity-based ones.

- The blue cluster from time, with an initial significant response but with the lagged response peak at 3 hours and 15 minutes, with the green cluster from popularity, the ones just below the most popular.

Figure 9 shows a gross-granularity association between temporal clusters and the popularity-based ones, to give an intuition on the best possible mapping between the two aspects. However, notice that in summary, time-based and popularity-based clustering results show low correlation. This highlights the fact that popularity alone is not a sufficient analysis dimension for guessing/understanding the social response in time of given brands.

5 Prediction

Once we realized that simple popularity metrics cannot be good proxies for describing the temporal dynamics of a brand, we focused on determining the best prediction model for estimating the membership of a brand in one of the four clusters determined before. In practice, we consider as labels the set of 4 clusters and we aim at predicting the correct label for a given brand B .

5.1 Pre-processing phase

For building the prediction, we consider only the 68 brands featuring one event during the MFW, to avoid aliasing due to the overlapping impact of multiple events. Subsequently, in the test phase we check the results of the prediction on all the 73 brands.

Some pre-processing of the data was performed in order to convert categorical variables into numerical ones. This avoids problems with the models that requires non-categorical features only. Other transformations have been applied to: (1) normalize continuous variables to a common range; and (2) convert variables with types allowing for many representations (such as date and time) to a unique, common format.

5.2 Prediction Models

We define some trivial baseline predictors and then we compare them with more advanced ones. In the following we introduce each of them and we report on the method applied and on the results obtained.

Baselines. We define three baseline models (simple classifiers) against which we test the results of more advanced ones:

1. **Random** strategy: it predicts the membership with a uniform random distribution to all the possible labels;
2. **Most frequent** strategy: it always predicts the membership to the most frequent label in the training set;
3. **Stratified** strategy: it generates label predictions according to the probability distribution of the labels in the training set.

Naive Bayes classifier for multivariate Bernoulli models. This Naive Bayes classifier applies to data that is distributed according to multivariate Bernoulli distributions; i.e., we assume each feature to be a binary-valued variable. Therefore, samples are represented as binary-valued feature vectors. Variables can be natively binary or transformed into binary.

Logistic Regression. In regularized logistic regression, the probabilities describing the possible outcomes of a single trial are modeled using a logistic function. We implement a version that can fit a multiclass (with one-vs-rest paradigm) norm-2 penalized logistic regression with optional regularization C . This can be represented as an optimization problem that minimizes the following cost function:

$$\min_{w,C} \frac{1}{2} w^T w + C \sum_{i=1}^m \log(\exp(-y_i(\mathbf{X}_i^T w + c)) + 1). \quad (2)$$

Cross-Validated Logistic Regression. We also improve the basic Logistic Regression model performance through cross-validation. The function is given 10 values in log-scale between 0.0001 and 10000 for the parameter C , the best hyperparameter is selected by the Stratified K-fold cross-validator. Being a multiclass problem, we compute the hyperparameters for each class using the best scores got by doing a one-vs-rest in parallel across all folds and classes.

Support Vector Machine. Support vector machines are a set of supervised learning methods used for classification, regression and outlier detection.

We adopt a One-versus-rest strategy, thus the problem can be formulated at each iteration as a 2-class problem. Given training vectors $x_i \in R^p$, with

$i = 1, \dots, n$, in two classes, and a vector $y \in \{1, -1\}^n$, the Support Vector Classifier solves the following dual minimization problem:

$$\min_a \frac{1}{2} a^T \mathbf{Q} a - e^T a \quad (3)$$

subject to:

$$y^T a = 0, 0 \leq a_i \leq C, i = 1, \dots, n. \quad (4)$$

where e is the vector of all ones, $C > 0$ is the upper bound, \mathbf{Q} is an $n \times n$ positive semidefinite matrix, $\mathbf{Q}_{ij} = y_i y_j K(x_i, x_j)$, where $K(x_i, x_j) = \phi(x_i)^T \phi(x_j)$ is the kernel function. Training vectors are implicitly mapped into a higher dimensional space by the function ϕ . The decision function applied by SVM is therefore:

$$\text{sign}\left(\sum_{i=1}^n y_i a_i K(x_i, x) + b\right) \quad (5)$$

Decision Tree. Decision Trees are a non-parametric supervised learning method used for classification and regression. The goal is to create a model that predicts the value of a target variable by learning simple decision rules inferred from the data features.

Random Forest This is a perturb-and-combine technique specifically designed for trees. This means a diverse set of classifiers is created by introducing randomness in the classifier construction. The prediction of the ensemble is given as the averaged prediction of the individual classifiers. With random forest, we build each tree in the ensemble starting from a sample, drawn with replacement (i.e., a bootstrap sample) from the training set. In addition, when splitting a node during the construction of the tree, the split that is chosen is no longer the best split among all features. Instead, the split that is picked is the best split among a random subset of the features, hence yielding an overall better model.

5.3 Fitting the models

We fit the models described in the previous section, in order to compare the results and pick the best one in the final classification step. The different models were fit with different features in order to find the best combination of input variables for the final classification and probability estimation. The final selected features are: $[x^{\text{start_minute}}, x^{\text{start_hour}}, x^{\text{start_day}}, x^{\text{end_minute}}, x^{\text{end_hour}}, x^{\text{end_day}}, x^{\text{class}}, x^{\text{live}}, x^{\text{type}}, x^{\text{invitation}}, x^{\text{open}}]$.

Table 2 reports the main performance coefficients of the models, covering: time performance indicators (time needed for performing the fitting and the prediction), weak indicators, and cross-validation indicators. Indeed, we applied Leave-One-Out Cross Validation (LOO-CV) for all the adopted models, in order to obtain better performance for the different models. This has been possible thanks to the relatively small size of the data-set.

More precisely, the columns in Table 2 describe:

1. the classifier adopted for the problem;
2. the time spent during the fitting phase, when training the model;
3. the time spent for the final prediction, for classifying the training samples;
4. the number of misclassified elements in the training (error);
5. the log loss on the training;
6. the number of misclassified elements in the test adopting LOO-CV (error);
7. the average log loss on the test adopting LOO-CV;
8. the standard deviation of the log loss on the test adopting LOO-CV;
9. the average log loss on the training adopting LOO-CV;
10. the standard deviation of the log loss on the training LOO-CV.

The number of errors for the random strategy is an average of the number of errors above 100 runs, in order to keep this measure more confident.

5.4 Results and Discussion

Execution time. In terms of execution time, all these models are fast both in the training phase and in the prediction phase. This is also related to the fact that the dataset is relatively small. The only two techniques that are slower than the others are Cross-Validated Logistic Regression (due to the cross-validation phase that has to choose the best value for the parameter C) and Random Forest (because of the diverse set of classifiers created by introducing randomness in their construction, and because of the high number of different estimators which we set to 1000).

Evaluation of training. Regarding the number of errors obtained when using as test set the training self itself, it is well known that this type of predictions leads to some significant overfitting; however, with small data-sets like ours, also this measure can be considered. For this performance indicator we have a predominance of the tree-based algorithms, i.e., Decision Trees and Random Forest. However, the respective performance accuracy is a clear example of overfitting: indeed, decision-tree learners can often create over-complex trees that do not generalise the data well. In any case, each method is outperforming the baselines Dummy Most Frequent and Dummy Stratified.

From the *log loss test* columns, where again test and train set are perfectly overlapping, we can notice the overfitting of the tree-based learners, together with nice performance in probability estimation from the other models, too.

Evaluation of LOO-CV. Regarding the *Leave One Out Cross-Validation* indicators, as we said we have applied prediction one sample at a time, using all the remaining ones as training set. This avoids the overfitting problem and thus are more accurate and truthful performance indicators. As one can expect, the number of errors are increasing. For this indicator, the best model is the Random Forest, that seems to beat the overfitting problem of the decision trees, upon which it is based. Every model we adopted is outperforming the baseline methods.

Table 2. Performance of all the techniques adopted in the classification problem. From left to right: the classifier adopted, the time spent during the fitting phase, the time spent for the final prediction, the number of training errors, the log loss on the training, the number of errors by adopting leave-one-out cross validation, the average log loss on the test (one row at a time) with the standard error adopting leave-one-out cross validation, the average log loss on the training with standard error adopting leave-one-out cross validation.

classifier	<i>Time indicators (s.)</i>		<i>Weak indicators</i>		<i>Leave-One-Out Cross-Validation indicators</i>				
	Fitting	Prediction	Errors	Log loss	Errors	Log loss test		Log loss training	
						Mean	Std	Mean	Std
Bernoulli NaiveBayes	0,0035	0,0015	22	1,39806	36	2,76323	4,17595	1,3901	0,0349
Logistic Regression	0,0225	0,0015	17	0,61651	33	0,76421	0,72045	0,6129	0,0123
Cross-Validated Log. Regr.	2,7521	0,0015	21	0,86406	32	0,56453	0,36066	0,8162	0,0666
Support Vector Machine	0,0190	0,0060	15	1,03000	38	0,49924	0,24064	0,9835	0,0444
Decision Tree	0,0035	0,0020	2	0,03798	37	17,03322	17,26817	0,0375	0,0044
Random Forest	1,1363	0,2067	2	0,25385	28	0,83452	0,79575	0,2507	0,0052
Dummy MostFrequent	0,0020	0,0020	36	17,03282	36	17,50636	17,26817	17,0328	0,2398
Dummy Stratified	0,0020	0,0015	48	20,81789	48	13,24806	16,79498	21,6853	1,8348
Uniform Distribution		0,0005		1,38629					
Random Distribution		0,0020		1,75860	60,47000			1,7586	0,0980

Considering the probability estimation indicators, i.e., log loss, we immediately notice the overfitting of the Decision Tree, as expected from the previous analysis. Indeed, the average log loss on the test for this type of estimator is really close to the baselines strategies Most Frequent and Stratified. Together with the test log loss standard deviation, these measures are proving the high variance of the over-complex model built by the tree. The best model considering the log loss on the test samples (both in terms of mean and standard deviation) is therefore the Support Vector Machine, with the lowest mean of the log loss and also the lowest standard deviation. The remaining indicators of the log loss on the training sets underline once again the overfitting issue of the Decision Tree and fairly good performance of the other models.

Summary. Concluding, if we give more importance to the Leave-One-Out Cross-Validation indicators, there are two models that can be preferred to the others: the first one is the Random Forest, which has a really low number of errors (only 2 over 68) when considering all the data-set as training and test set, and maintains the lowest number of errors (28 over 68) among the other techniques in the classification phase of the Cross-Validation. The second is the Support Vector Machine, which has the best performance in probability estimation, as we can see from the log loss LOO-CV indicators. However, this technique reports a significant number of errors (38 errors) in the classification phase of the Cross-Validation.

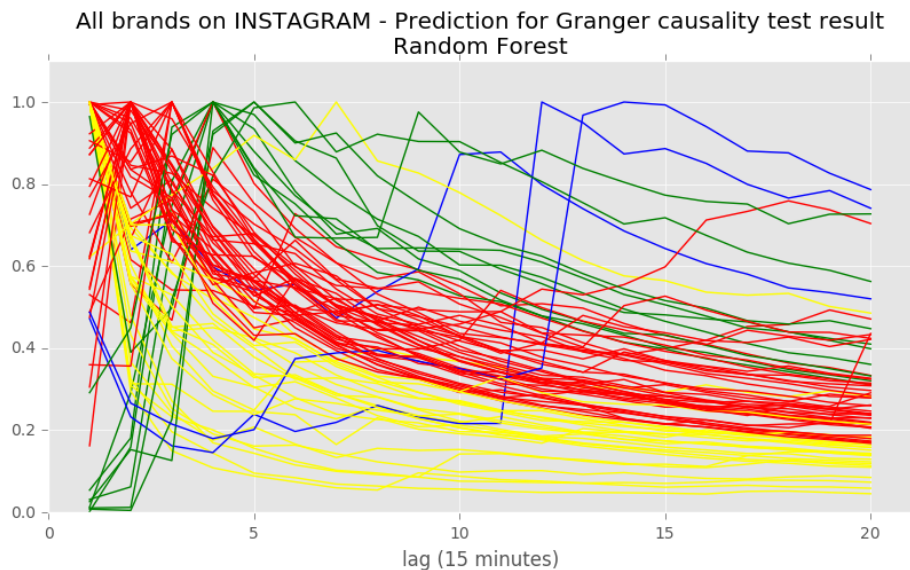


Fig. 10. Clusters produced by the Random Forest predictor in terms of Granger causality curves.

6 Related work

Various studies analyzed the impact of events on social media. We consider both works that deal with response to events on the social media in general, and works that focus specifically on the role of social media in the fashion domain.

Social media event response. The work [10] is concerned with 21 hot events which were widely discussed on Sina Weibo; it empirically analyzes their posting and reposting characteristics. In the work [2], by automatically identifying events and their associated user-contributed social media documents, the authors show how they can enable event browsing and search in a search engine. The work [3] underlines how user-contributed messages on social media sites such as Twitter have emerged as powerful, real-time means of information sharing on the Web. The authors explore approaches for analyzing the stream of Twitter messages to distinguish between messages about real-world events and non-event messages.

The focus of [7] is to detect events and related photos from Flickr, by exploiting the tags supplied by users. This task is hard because (1) Flickr data is noisy; (2) capturing the content of photos is not easy. They distinguish between tags of aperiodic events and those of periodic events; for both classes, event-related tags are clustered such that each cluster consists of tags with similar temporal and locational distribution patterns as well as with similar associated photos. Finally, for each tag cluster, photos corresponding to the represented event are extracted.

The problem of event summarization using tweets is well faced by [6], where the authors argue that for some highly structured and recurring events, such as sports, it is better to use sophisticated techniques to summarize the relevant tweets via Hidden Markov Models.

The paper [5], adding the information given by cell-phone traces, deals with the analysis of crowd mobility during special events. They analyze nearly 1 million cell-phone traces and associate their destinations with social events. They show that the origins of people attending an event are strongly correlated to the type of event.

Finally, [1] proposes a procedure consisting of a first collection phase of social network messages, a subsequent user query selection, and finally a clustering phase, for performing a geographic and temporal exploration of a collection of items, in order to reveal and map their latent spatio-temporal structure. Specifically, both several geo-temporal distance measures and a density-based geo-temporal clustering algorithm are proposed. The paper aims at discovering the spatio-temporal periodic and non-periodic characteristics of events occurring in specific geographic areas.

Social Media Analysis for Fashion. The work [12] presents a qualitative analysis on the influence of social media platforms on different behaviors of fashion brand marketing. They analyze their styles and strategies of advertisement. The authors employ both linguistic and computer vision techniques. The study

[11] set out to identify attributes of social media marketing (SMM) activities and examine the relationships among those perceived activities, value equity, relationship equity, brand equity, customer equity, and purchase intention through a structural equation model.

The findings of [14] show that different drivers influence the number of likes and the number of comments to fashion posts. Namely, vivid and interactive brand post characteristics enhance the number of likes. An analysis in [8] was conducted during a 2011 Victoria's Secret Fashion Show that reference the show. Although the majority were idiosyncratic remarks, many tweets contain evidence of social status comparisons to the fashion models.

The article [9], based on two studies of the fashion industry, examines one of its key institutions, London Fashion Week (LFW). The authors argue that this event is a materialization of the field of fashion. They examine how LFW renders visible the boundaries, relational positions, capital and habitus at play in the field, reproducing critical divisions within it.

Finally, [13] develops a motion capture system using two cameras that is capable of estimating a constrained set of human postures in real time. They first obtain a 3D shape model of a person to be tracked and create a posture dictionary consisting of many posture examples.

7 Conclusions

We discussed how social content spreads in time as a response to live events, focusing on the Milano Fashion Week. We demonstrated that brands can be clustered into 4 classes of increasingly prolonged responses, from ones showing a peak which is close in time to the event, to ones which do not show sharp peaks but rather a smoothed behavior; we also showed that brand popularity alone is not sufficient for explaining such a difference. Estimating the time latency of social responses to events is important for the brands, which could plan reinforcement actions more accurately, essentially by adding few but well-planned social actions so as to sustain their social presence over time. Our future work is to attempt to correlate spreading in time with other features beyond brand popularity, e.g. by studying the profiles of each brand's social networks and specifically of Instagram.

8 Acknowledgments

We thank the FashionInProcess group (<http://www.fashioninprocess.com/the-collective>) of Politecnico di Milano, and especially Paola Bertola, Chiara Colombi and Federica Vacca, who supported us in the definition of the domain-specific knowledge related to the event.

References

1. Arcaini, P., Bordogna, G., Ienco, D., Sterlacchini, S.: User-driven geo-temporal density-based exploration of periodic and not periodic events reported in social networks. *Information Sciences* pp. 122–143 (2016)
2. Becker, H., Naaman, M., Gravano, L.: Learning Similarity Metrics for Event Identification in Social Media (2010)
3. Becker, H., Naaman, M., Gravano, L.: Beyond Trending Topics: Real-World Event Identification on Twitter (2011)
4. Bradford, J.: *Fashion Journalism* (2014)
5. Calabrese, F., Pereira, F., Di Lorenzo, G., Liu, L., Ratt, C.: The Geography of Taste: Analyzing Cell-Phone Mobility and Social Events. *Pervasive 2010* pp. 22–37 (2010)
6. Chakrabarti, D., Punera, K.: Event Summarization using Tweets. *Yahoo! Research* (2011)
7. Chen, L., Roy, A.: Event Detection from Flickr Data through Wavelet-based Spatial Analysis (2009)
8. Chrisler, J., Fung, K., Lopez, A., Gorman, J.: Suffering by comparison: Twitter users' reactions to the Victoria's Secret Fashion Show. *Body Image* pp. 648–652 (2013)
9. Entwistle, J., Rocamora1, A.: The Field of Fashion Materialized: A Study of London Fashion Week. *Journal of Sociology* pp. 735–751 (2006)
10. Guan, W., Gao, H., Yang, M., Li, Y., Ma, H., Qian, W., Cao, Z., Yang, X.: Analyzing user behavior of the micro-blogging website Sina Weibo during hot social events. *Physica A* pp. 340–351 (2013)
11. Kim, A., Ko, E.: Do social media marketing activities enhance customer equity? An empirical study of luxury fashion brand. *Journal of Business Research* (2011)
12. Manikonda, L., Venkatesan, R., Kambhampati, S., Li, B.: Trending Chic: Analyzing the Influence of Social Media on Fashion Brands. Department of Computer Science, Arizona State University (2016)
13. Okada, R., Stenger, B., Ike, T., Kondo, N.: Virtual Fashion Show Using Real-Time Markerless Motion Capture. Corporate Research & Development Center, Toshiba Corporation (2006)
14. de Vries, L., Gensler, S., Leeflang, P.: Popularity of Brand Posts on Brand Fan Pages: An Investigation of the Effects of Social Media Marketing. *Journal of Interactive Marketing* (2012)