

Towards Miss Universe Automatic Prediction: The Evening Gown Competition

Johanna Carvajal^{†◊}, Arnold Wiliem[†], Conrad Sanderson^{†◊}, Brian Lovell[†]

[†] University of Queensland, Brisbane, Australia

[◊] Data61, CSIRO, Australia

Abstract—Can we predict the winner of Miss Universe after watching how they stride down the catwalk during the evening gown competition? Fashion gurus say they can! In our work, we study this question from the perspective of computer vision. In particular, we want to understand whether existing computer vision approaches can be used to automatically extract the qualities exhibited by the Miss Universe winners during their catwalk. This study can pave the way towards new vision-based applications for the fashion industry. To this end, we propose a novel video dataset, called the Miss Universe dataset, comprising 10 years of the evening gown competition selected between 1996-2010. We further propose two ranking-related problems: (1) Miss Universe Listwise Ranking and (2) Miss Universe Pairwise Ranking. In addition, we also develop an approach that simultaneously addresses the two proposed problems. To describe the videos we employ the recently proposed Stacked Fisher Vectors in conjunction with robust local spatio-temporal features. From our evaluation we found that although the addressed problems are extremely challenging, the proposed system is able to rank the winner in the top 3 best predicted scores for 5 out of 10 Miss Universe competitions.

I. INTRODUCTION

Miss Universe is a worldwide pageant competition held every year since 1952 and is organised by *The Miss Universe Organization* [2]. Every year up to 89 candidates participate in the competition. Each delegate must first win their respective national pageants. Miss Universe is broadcast in more than 190 countries around the world and is watched by more than half a billion people annually [2], [3]. The format has slightly changed during the 64 year period. However, the most common competition format is as follows. All candidates are preliminary judged in three areas of competition: Interview, Swimsuits and Evening Gown. After that, the top 10 or 15 semi-finalists are short-listed during the coronation night. The semi-finalists compete again in swimsuits and evening gowns. The best 5 finalists are selected and go through an interview round. Finally, the runners-up and winner are announced.

Although Miss Universe is one of the most publicised beauty pageants in the world, it is not the only existing pageant competition. A list of beauty pageants from around the world includes up to 22 events among international, continental and, regional pageants. Moreover, there are more than 260 national pageants. In the US alone, there are approximately 28 national pageants [1].

During the swimsuit and evening gown competition, the catwalk is judged by several aspects. Candidates must emanate poise, posture, grace, elegance, balance, confidence, energy, charisma, and sophistication. Additionally, during the swimsuit competition candidates are expected to have a well-proportioned body, good muscle tone, proper level of body fat and show fitness and body shape. In our work, we aim to capture these qualities to predict the winner. This can pave the way of numerous vision-based applications for the fashion industry such as automatic training systems for amateur models who aspire to become professionals. Due to the complexity

of this problem, we propose to initially study the evening gown competition. To this end, we collect a new dataset of videos recorded during the evening gown competition where the judges' scores are publicly available.

As mentioned, there are many potential commercial application for an automatic system able to analyse and predict the best catwalk in a beauty pageant. Automatically predicting the winner can be useful for specialised betting sites such as Odds Shark, Sports Bet, Bovada, and Bet Online. These betting sites allow the audience to bet for their favourite candidate in Miss Universe. Moreover, the catwalk analysis can be a powerful tool for boutique talent agencies such as *Polished by Donna* that provides training for improving the catwalk and offer their services to future beauty pageants candidates [4]. For boutique talent agencies, an automatic catwalk analysis system can help to compare the catwalk of each amateur model against herself or against an experienced catwalker.

Towards automatic prediction of Miss Universe, we first collect a novel Miss Universe (MU) dataset. The dataset comprises 10 years of Miss Universe selected from 1996 to 2010. Only those years with available videos and official scores are selected. The years included in this datasets are: 1996, 1997, 1998, 1999, 2000, 2001, 2002, 2003, 2007, and 2010. The years not included were due to the videos and/or the scores were not publicly available.

It comprises 105 videos and 18,343 frames depicting each candidate catwalk in the evening gown competition. Fig. 1 shows two examples of best and worst judges' scores during the evening gown competition. We propose two sub-problems: (1) Miss Universe Listwise Ranking (MULR), and (2) Miss Universe Pairwise Ranking (MUPR). The MULR problem aims to predict the winner of the evening gown competition, which can be useful during a beauty pageant competition and also for betting sites. The MUPR problem focuses on judging the catwalk between two participants or to see the improvement of one model's catwalk. The solution of the MUPR problem could be used for developing applications for boutique talent agencies.

In this work, we propose an approach which will address both problems simultaneously. More specifically, we found that it is possible to share the model trained from one problem with the other problem. We use our approach in conjunction with the video descriptors used for action analysis in [11]. In particular the video descriptors are extracted on a pixel-base and make use of gradients and optical flow. Gradients and optical flow have been shown to be effective for video representation. Then, the video descriptors are encoded using the Stacked Fisher Vectors (SFV) approach, which has recently shown successful performance for action analysis [26]. From our evaluations, we found that that our proposed problems are extremely challenging. However, further analysis suggests that both problems could still be potentially solved using a computer vision approach.

Contributions — we present 4 main contributions: (1) we study two novel problems for automatic ranking of Miss Universe evening gown competition participants using computer vision techniques; (2) we propose a novel dataset called the Miss Universe (MU) dataset that comprises 10 years of the Miss Universe evening gown competition selected between 1996-2010; (3) we propose an approach that addresses both problems simultaneously; (4) we adapt recent video descriptors, shown to be effective in action analysis, into our framework.

We continue our paper as follows. Section II summarises the related work. The two sub-problems for automatically predicting the winner of Miss Universe during the evening gown competition are explained in Section III. In Section IV we present our proposed approach that simultaneously addresses the two problems. Section V describes the Miss Universe dataset, the evaluation protocol, and the evaluation metrics. In Section VI we present the results for both sub-problems. The main findings are summarised in Section VII.

II. RELATED WORK

An automatic system to predict the best catwalk in a beauty pageant has not been investigated before. The catwalk during the evening gown competition can be seen as a walk assessment, action assessment, or fine-grained action analysis. For catwalk assessment in Miss Universe, judges assess the quality of the walk. Catwalk analysis can be also related to fine-grained action analysis, where the aim is to distinguish the fine and subtle differences between two candidate catwalks. In the following sections, we summarise recent works for catwalk Assessment, Fine-Grained Action Analysis and some popular features for action analysis.

Action Assessment — Gait and walk assessments have been investigated for elderly people and humans with neurological disorders [34], [16]. Two web-cams are used to extract gait parameters including walking speed, step time, and step length in [34]. The gait parameters are used for a fall risk assessment tool for home monitoring of older adults. For rehabilitation and treatment of patients with neurological disorders, automatic gait analysis with a Microsoft Kinect sensor is used to quantify the gait abnormality of patients with multiple sclerosis [16]. A gait analysis system consisting of two camcoders located on the right and left side of a treadmill is employed in [25]. This system fully reconstructs the skeleton model and demonstrates good accuracy compared to Kinect sensors. Despite being a related problem, for our Miss Universe catwalk analysis, Kinect sensors or multi-cameras are simply not available. The assessment of quality of actions using only visual information is still under early development. A recent work to predict the expert judges' scores for actions diving and figure skating in the Olympic games is presented in [28]. The concept behind the score prediction is to learn how to assess the quality of actions in videos.

Fine-Grained Action Analysis — Catwalk analysis can be also related to fine-grained action analysis. Fine-grained action analysis has been recently investigated for action recognition [10], [14], [22], [29], [31], where it is important to recognise small differences in activities such as cut and peel in food preparation. This is in contrast to traditional action

recognition where the goal is to recognise full-body activities such as walking or jumping.

Features for Action Analysis — Improved dense trajectory (IDT) features in conjunction with Fisher Vector representation have recently show outstanding performance for the action recognition problem [35]. This approach densely samples feature points at several spatial scales in each frame and tracks them using optical flow. For each trajectory the following descriptors are computed: Trajectory, Histogram of Gradients, Histogram of Optical Flow, and Motion Boundary histogram. Finally, all descriptors are concatenated and normalised. IDT features are also popular for fine-grained action recognition [22], [29], [31]. However, some disadvantages have been reported. IDT generates irrelevant trajectories that are eventually are discarded. Processing such trajectories is time consuming and hence not suitable for realistic environments [5], [19].

Gradients have been used as a relatively simple yet effective video representation [11]. Each pixel in the gradient image helps extract relevant information, eg. edges of a subject. Gradients can be computed at every spatio-temporal location (x, y, t) in any direction in a video. Lastly, since the task of action recognition is based on an ordered sequence of frames, optical flow can be used to provide an efficient way of capturing local dynamics and motion patterns in a scene [17].

III. PROBLEM DEFINITION

During the evening gown competition, candidates are given an average score based on their catwalk. Different judges are selected each year to score each candidate. This score is used in conjunction with the swimming competition, to select the best 5 finalists, where finally the Miss Universe winner is announced. Candidates with the best scores strut with attitude down the catwalk projecting confidence. See top row of Fig. 1 for examples. Their arms are kept relaxed and swing naturally with the body. In general, they exhibit a flouncing walk and ooze elegance as they stalk the runway. Candidates with the worst scores tend to exhibit issues such as stiff arms (resulting in robotic or awkward appearance) and drooping their heads. See bottom row of Fig. 1 for examples. It can be also seen that the candidate with the worst catwalk during Miss Universe 2010 (Fig. 1 bottom left) finds herself struggling to walk with the ribbon dress that is too tight for her.

Our central problem is to predict the best catwalk during the evening gown competition. This can be considered as an instance of the ranking problem. The ranking problem has been explored in various domains such as collaborative filtering, documents retrieval, and sentiment analysis [9].

In our work, we define two ranking sub-problems: (1) Miss Universe Listwise Ranking (MULR), and (2) Miss Universe Pairwise Ranking (MUPR). While MULR focuses on rank ordering of all Miss Universe participants in the same year, MUPR considers pairwise comparisons of two participants in the same year. We note that these two sub-problems have also been described in [12], [13] for general machine learning problems.

A. Miss Universe Listwise Ranking (MULR) Problem

The MULR problem can be formalised as follows. Given a query $Q_i = \{\mathbf{p}_j^{(q)}\}_{j=1}^{N_q}$, where $\mathbf{p}_j^{(q)}$ is the video of a participant for Miss Universe from year q and N_q is the total number



Fig. 1: Examples of best and worst scores for Miss Universe versions 2003 and 2010

of candidates for that specific year. Let $\mathcal{G}_l = \{\mathcal{S}_m\}_{m=1}^M$ be the gallery containing M sets of Miss Universe from M years, where $\mathcal{S}_m = \{\mathbf{p}_j^{(m)}\}_{j=1}^{N_m}$ is the set of N_m participant videos of Miss Universe from year m . Each set of participants \mathcal{S}_m is associated with a set of judgements (scores) $\mathbf{y}^{(m)} = [y_1^{(m)}, \dots, y_{N_m}^{(m)}]$. The judgement $y_j^{(m)}$ represents the average score of participant $\mathbf{p}_j^{(m)}$. We note that the average score is calculated by averaging the scores given by all the judges during the evening gown competition. A set of video descriptors $\mathbf{v}_j^{(m)} = \Phi(\mathbf{p}_j^{(m)})$, where $\mathbf{v}_j^{(m)} \in \mathbb{R}^d$ are extracted from each participant video, $\mathbf{p}_j^{(m)}$.

Let $f_l : \mathbb{R}^d \mapsto \mathbb{R}^1$ be a scoring function that calculates a participant score based on its corresponding video descriptors. Given the query \mathcal{Q}_l , the function f_l can automatically score each participant in \mathcal{Q}_l . Let $\mathbf{y}^{(q)} = [y_1^{(q)}, \dots, y_{N_q}^{(q)}]$ be the actual score from the judges for participants in the query \mathcal{Q}_l , and $\hat{\mathbf{y}}^{(q)} = [f(y_1^{(q)}), \dots, f(y_{N_q}^{(q)})]$ be the estimated score of function f trained using the gallery set \mathcal{G}_l , the main task in MULR problem is to find the best f_l , where ideally the ranking of $\hat{\mathbf{y}}^{(q)}$ is the same as $\mathbf{y}^{(q)}$.

B. Miss Universe Pairwise Ranking (MUPR) Problem

For the MUPR problem, we first consider a gallery, $\mathcal{G}_p = \{(\mathbf{p}_l^{(m)}, \mathbf{p}_k^{(m)})\}_{m=1}^M, \mathbf{p}_l^{(m)}, \mathbf{p}_k^{(m)} \in \mathcal{S}_m, l \neq k$ wherein each element in the gallery is a pair of participant videos from the same year of Miss Universe. Note that the gallery \mathcal{G}_p considered in this problem is different from the gallery \mathcal{G}_l considered in MULR problem. Each pair in the gallery has its corresponding label $y_{lk}^{(m)}$ which is defined via:

$$y_{lk}^{(m)} = \begin{cases} +1; & y_l^{(m)} > y_k^{(m)} \\ -1, & \text{otherwise} \end{cases}, \quad (1)$$

where $y_l^{(m)}$ and $y_k^{(m)}$ are the actual score from the judges. Let $(\mathbf{p}_l^{(q)}, \mathbf{p}_k^{(q)})$, $y_{lk}^{(q)}$ be a query pair and its corresponding label, the main task for the MUPR problem is to find the best ranking function $f_p(\cdot) = \{-1, +1\}$ where ideally $y_{lk}^{(q)} = f_p(\mathbf{p}_l^{(q)}, \mathbf{p}_k^{(q)})$.

IV. PROPOSED APPROACH

We first describe the video descriptors used in our work. We then present our approach to solve both MULR and MUPR problems simultaneously.

A. Video Descriptors

Here, we describe how to extract from a video a set of features on a pixel level. A video $\mathcal{V} = \{\mathbf{I}_t\}_{t=1}^T$ is an ordered set of T frames. Each frame $\mathbf{I}_t \in \mathbb{R}^{r \times c}$ can be represented by a set of feature vectors $\mathbf{F}_t = \{\mathbf{f}_p\}_{p=1}^{N_t}$. We extract the following $d = 14$ dimensional feature vector for each pixel in a given frame t [11]:

$$\mathbf{f} = [x, y, \mathbf{g}, \mathbf{o}]^\top \quad (2)$$

where x and y are the pixel coordinates, while \mathbf{g} and \mathbf{o} are:

$$\mathbf{g} = \left[|J_x|, |J_y|, |J_{yy}|, |J_{xx}|, \sqrt{J_x^2 + J_y^2}, \text{atan} \frac{|J_y|}{|J_x|} \right] \quad (3)$$

$$\mathbf{o} = \left[u, v, \frac{\partial u}{\partial t}, \frac{\partial v}{\partial t}, \left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right), \left(\frac{\partial v}{\partial x} - \frac{\partial u}{\partial y} \right) \right] \quad (4)$$

The first four gradient-based features in Eq. (3) represent the first and second order intensity gradients at pixel location (x, y) . The last two gradient features represent gradient magnitude and gradient orientation. The optical flow based features in Eq. (4) represent: the horizontal and vertical components of the flow vector, the first order derivatives with respect to time, the divergence and vorticity of optical flow [6], respectively. With this set of descriptors, we aim to capture the following attributes: shape with the coordinates, appearance with the gradients, and motion with the optical flow. Typically only a subset of the pixels in a frame correspond to the object of interest ($N_t < r \times c$). As such, we are only interested in pixels with a gradient magnitude greater than a threshold β [18]. We discard feature vectors from locations with a small magnitude, resulting in a variable number of feature vectors per frame. For each video \mathcal{V} , the feature vectors are pooled into set $\mathcal{F} = \{\mathbf{f}_n\}_{n=1}^N$ containing N vectors.

B. Stacked Fisher Vectors

The traditional Fisher Vector (FV) consists in describing a pooled set of features by its deviation from a generative model. FV encodes the deviations from a probabilistic version of a visual dictionary, which is typically a Gaussian Mixture Model (GMM) with diagonal covariance matrices [27], [32]. The parameters of a GMM with K components can be expressed as $\lambda = \{w_k, \boldsymbol{\mu}_k, \boldsymbol{\sigma}_k\}_{k=1}^K$, where, w_k is the weight, $\boldsymbol{\mu}_k$ is the mean vector, and $\boldsymbol{\sigma}_k$ is the diagonal covariance matrix for the k -th Gaussian. The parameters are learned using the Expectation Maximisation algorithm [8] on training data. Given the pooled set of features \mathcal{F} from video \mathcal{V} , the deviations from the GMM are then accumulated using [32]:

$$\mathcal{G}_{\boldsymbol{\mu}_k}^{\mathcal{F}} = \frac{1}{N\sqrt{w_k}} \sum_{n=1}^N \gamma_n(k) \left(\frac{\mathbf{f}_n - \boldsymbol{\mu}_k}{\boldsymbol{\sigma}_k} \right) \quad (5)$$

$$\mathcal{G}_{\boldsymbol{\sigma}_k}^{\mathcal{F}} = \frac{1}{N\sqrt{2w_k}} \sum_{n=1}^N \gamma_n(k) \left[\frac{(\mathbf{f}_n - \boldsymbol{\mu}_k)^2}{\boldsymbol{\sigma}_k^2} - 1 \right] \quad (6)$$

where vector division indicates element-wise division and $\gamma_n(k)$ is the posterior probability of \mathbf{f}_n for the k -th component:

$$\gamma_n(k) = \frac{w_k \mathcal{N}(\mathbf{f}_n | \boldsymbol{\mu}_k, \boldsymbol{\sigma}_k)}{\sum_{i=1}^K w_i \mathcal{N}(\mathbf{f}_n | \boldsymbol{\mu}_i, \boldsymbol{\sigma}_i)} \quad (7)$$

The Fisher vector for each video \mathcal{V} is represented as the concatenation of $\mathcal{G}_{\mu_k}^{\mathcal{F}}$ and $\mathcal{G}_{\sigma_k}^{\mathcal{F}}$ (for $k = 1, \dots, K$) into vector $\mathcal{G}_{\lambda}^{\mathcal{F}}$. As $\mathcal{G}_{\mu_k}^{\mathcal{F}}$ and $\mathcal{G}_{\sigma_k}^{\mathcal{F}}$ are d -dimensional, $\mathcal{G}_{\lambda}^{\mathcal{F}}$ has the dimensionality of $2dK$. Power normalisation is then applied to each dimension in $\mathcal{G}_{\lambda}^{\mathcal{F}}$. The power normalisation to improve the FV for classification was proposed in [27] of the form $z \leftarrow \text{sign}(z)|z|^{\rho}$, where z corresponds to each dimension and the power coefficient $\rho = 1/2$. Finally, l_2 -normalisation is applied. Note that we have omitted the deviations for the weights as they add little information [32].

Stacked Fisher Vectors (SFV) is a multi-layer representation of standard FV [26]. SFV first performs traditional FV representation over densely sampled subvolumes based on low level descriptors. The extracted FVs have a high dimensionality and are fed the next layer. The second layer reduces the obtained FVs, and then those reduced FVs are encoded again with FV representation.

C. Classification

We address both MULR and MUPR problems using the same framework. Recall that the main objective of the MULR problem is to find the best $f_l(\cdot)$ wherein its scores can be used to rank the Miss Universe participants from the same year. We model such a function as a linear regression function:

$$f_l(\mathbf{v}) = \mathbf{w}^{\top} \mathbf{v} + b, \quad (8)$$

where \mathbf{w} and b are the parameters of the regression model and $\mathbf{v} \in \mathbb{R}^d$ is the extracted video descriptor after applying SFV. As it is not trivial to train the regression given the gallery \mathcal{G}_l with its corresponding actual ranking, we solve this problem by addressing MUPR, which is a much easier problem. This is possible as the ranking function f_p can be defined in terms of the scoring function f_l :

$$f_p(\mathbf{v}_l, \mathbf{v}_k) = \text{sign}(f_l(\mathbf{v}_l) - f_l(\mathbf{v}_k)), \quad (9)$$

where $\text{sign}(\cdot)$ only takes the sign of the input. Plugging the scoring function model into the above equation we obtain:

$$\begin{aligned} f_p(\mathbf{v}_l, \mathbf{v}_k) = & \text{sign}(f_l(\mathbf{v}_l) - f_l(\mathbf{v}_k)) \\ & \text{sign}(\mathbf{w}^{\top} \mathbf{v}_l + b - \mathbf{w}^{\top} \mathbf{v}_k - b) \\ & \text{sign}(\mathbf{w}^{\top} (\mathbf{v}_l - \mathbf{v}_k)) \\ & \text{sign}(\mathbf{w}^{\top} \mathbf{z}), \end{aligned} \quad (10)$$

where $\mathbf{z} \in \mathbb{R}^d$ is the new descriptor extracted via: $\mathbf{z} = \mathbf{v}_l - \mathbf{v}_k$. Notice that both f_l and f_p share the same model parameter \mathbf{w} . As we only focus on the ranking for MULR problem, the bias parameter, b in Eq. (8) can be excluded; the regression model thus becomes:

$$f_l(\mathbf{v}) = \mathbf{w}^{\top} \mathbf{v}. \quad (11)$$

With the above modification, we only need to perform the training step once for both functions. To this end, we perform the training step for the ranking function, f_p . Following the training formulation from the RankSVM described in [21]:

$$\frac{1}{2} \|\mathbf{w}\|_2^2 + C \sum_{m=1}^M \sum_{\mathbf{v}_l^{(m)}, \mathbf{v}_k^{(m)} \in \mathcal{S}_m} \ell(y_{lk}^{(m)} \mathbf{w}^{\top} \mathbf{z}_{lk}^{(m)}), \quad (12)$$

where $\mathbf{z}_{lk}^{(m)} = \mathbf{v}_l^{(m)} - \mathbf{v}_k^{(m)}$, is the new descriptor as described above; $y_{lk}^{(m)}$ is the ground truth for the MUPR problem described in Eq. (1), C is a training parameter and $\ell(\cdot)$ is the hinge loss.

V. MISS UNIVERSE (MU) DATASET

In this work, we propose the Miss Universe Dataset to address our problems. In particular, we have collected a novel dataset of videos depicting the evening gown competition for 10 years of Miss Universe (MU). The videos span from 1996 to 2010, where the judges' scores are available. The videos were downloaded from YouTube and the scores were obtained from the videos themselves or Wikipedia. Fig. 3 shows examples of scores. While the scores taken from the videos include each individual score from judge, only the average is used (circled in yellow).

We have collected 105 videos, 18,343 frames in total, with an average of 175 per video. Each video shows a candidate during the evening gown competition. Additionally, we manually select the bounding box enclosing each participant.

It is noteworthy to mention that the proposed MU dataset is extremely challenging due to variations in capture conditions for each year: (1) catwalk stage; (2) illumination conditions; (3) cameras capturing the event. As for the variations in cameras capturing the event, for our purpose we opted to use only one camera view depicting the longest walk without interruptions. Fig. 2 shows the catwalk stage for each year in the MU dataset. The dataset is available from <http://www.itee.uq.edu.au/sas/datasets>

A. Evaluation Protocol

We use leave-one-year-out protocol as the evaluation protocol for both MUPR and MULR. In particular, for each training-test set, we consider all participants from one year as the testing and the rest as training. As the dataset covers ten years of Miss Universe videos, there are ten training-test sets. Once the results from all the ten training-test sets are determined, the performance of a method is reported as the average of these results.

The MULR problem evaluation metric — In the MULR problem we are interested in evaluating how similar is the ranking determined to the scoring function f_l from the actual ranking of each year. To this end, we use the Normalized Discount Cumulative Gain (NDCG) proposed to measure ranking quality of documents [24], [30]. NDCG is often used to measure of the efficacy of web search algorithms [24].



Fig. 2: Catwalk stages for all years

Delegate [hide]	Evening Gown	Swimsuit	Average
Russia	9.64 (1)	9.88 (1) (7)	9.760 (1)
Venezuela	8.83 (5)	9.29 (2)	9.060 (2)
China	9.15 (2)	8.88 (5)	9.015 (3)
Panama	8.92 (3)	8.79 (7)	8.855 (4)
South Africa	8.79 (6)	8.90 (4)	8.845 (5)
Germany	8.84 (4)	8.81 (6)	8.825 (6)
Cyprus	8.49 (8)	9.15 (3)	8.820 (7)
Albania	8.51 (7)	8.34 (8)	8.425 (8)
India	8.10 (10)	8.32 (9)	8.210 (9)
Canada	8.39 (9)	7.99 (10)	8.190 (10)



Fig. 3: Judges' scores. Left: from Wikipedia. Right: from the video.

To use this metric, we consider each candidate video as a “visual” document. Here the rating of each visual document corresponds to the rank of the participant. Thus, we rate each visual document/participant video by assigning values between 1 to 10 with 10 being the highest score and 1 for the lowest. These values are assigned according to their corresponding rank. For instance, we assign the participant having the highest score with value 10 and assign the runner up with value 9.

In the original formulation, the NDCG measures the ranking quality based on the top b rated documents [24]:

$$\text{NDCG}_{@b} = \text{DCG}_{@b} / \text{IDCG}_{@b}, \quad (13)$$

where DCG is the discounted cumulative gain at particular rank position b and is defined as:

$$\text{DCG}_{@b} = \sum_{j=1}^b \frac{(2^{r(j)} - 1)}{\log_2(\max(2, j))} \quad (14)$$

The rating of the j -th participant in the ranking list is given by $r(j)$ and $\text{IDCG}_{@b}$ is the ideal DCG at position b . Note that $b = 1, \dots, N_q$ with N_q being the length of the ordering. A perfect list gets a $\text{NDCG}_{@b}$ score of 1. For our case, we always set $b = N_q$. We report the average percentage $\text{NDCG}_{@N_q}$ over all partitions and refer to it as the NDCG.

The MUPR problem evaluation metric — For the MUPR problem, we use the modified Kendall’s \mathcal{K}_τ as a performance measure discussed in [20]. \mathcal{K}_τ is defined as the number C of concordant pairs and the number D of discordant pairs. A pair $(p_l^{(m)}, p_k^{(m)})$ with $l \neq k$ is concordant, if $\hat{y}_{lk}^{(m)} = y_{lk}^{(m)}$. It is discordant if they disagree. The sum of C and D must be $\binom{N_q}{2}$. Kendall’s \mathcal{K}_τ can be defined as:

$$\mathcal{K}_\tau = \frac{C - D}{C + D} = 1 - \frac{2D}{\binom{N_q}{2}} \quad (15)$$

VI. EXPERIMENTS

To the best of our knowledge, this is the first work to study catwalk analysis for Miss Universe. We used the new Miss Universe dataset containing 10 versions of Miss Universe. Miss Universe 2003 contains 15 participants. The remaining versions each contain 10 participants. We used the bounding box enclosing the participant provided with the dataset. We resized all bounding boxes to 100×50 .

Setup — All videos were converted into gray-scale. We use the leave-one-year-out protocol, where we leave one version of Miss Universe out for testing. For each video, we extract a set of $d = 14$ dimensional features as explained in Section IV-A. Based on [11], we used $\beta = 40$, where β is the threshold used for selecting interesting low-level features. Parameters for the visual vocabulary GMM were learned using a large set of descriptors randomly obtained from training videos using the iterative Expectation-Maximisation algorithm [8]. The systems were implemented with the aid of the Armadillo C++ library [33]. Experiments were performed with three separate GMMs with varying number of components $K = \{256, 512, 1024\}$.

For the traditional FV representation, each video is represented by a FV. The FVs are fed to a linear SVM for classification.

For the first layer of SFV, we obtained a varying number of vectors using the traditional FV representation. Each vector is obtained using the low-level descriptors of 5 consecutive

TABLE I: Results for MUPR using \mathcal{K}_τ

Method	Visual Vocabulary Size		
	256	512	1024
FV (baseline)	52.63 %	52.16 %	52.03 %
SFV-PCA	56.73 %	51.81 %	50.98 %
SFV-RP	53.68 %	46.92 %	47.87 %

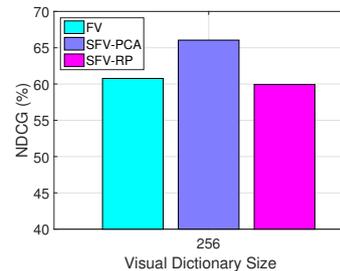


Fig. 4: Results for MULR using NDCG

frames. Then, we advanced by a frame and obtained a new FV. For the second layer of SFV, we reduced the dimensionality of each vector from layer 1 using two methods: Principal Component Analysis (PCA) and Random Projection (RP). For PCA, we retained the 90% of the energy [7]. For RP, we used the resulting dimensionality number obtained by PCA. We referred to these methods as SFV-PCA and SFV-RP.

Our classification model is described in Section IV-C. As explained, we address both problems using the same framework. We solve MULR by addressing MUPR first. In our implementation, we solve MUPR by using the LibLinear package [15] and set the bias parameter b to 0.

Results for MUPR — In Table I, we present the results for MUPR. The evaluation metric employed is Kendall’s \mathcal{K}_τ as per Eq. (15). From this table, we can see that our classification models using both dimensionality reduction techniques outperform the baseline FV representation. Using SFV-PCA with a visual dictionary size of 256 Gaussians leads to the best performance of 56.73%, which is 3.05 points higher than SFV-RP. PCA is an essential step for dimensionality reduction for this application. Despite the simplicity of random projection, its performance is inferior to PCA.

Results for MULR — Using the best setting for MUPR obtained with a visual vocabulary size of 256 Gaussians, we evaluated MULR. The evaluation metric employed is NDCG as per Eq. (13). Fig. 4 shows that SFV-PCA attained the best performance with 66.05%. Table II shows the individual performance using NDCG for each of the ten training-test sets as explained. Our SFV-PCA classification approach shows a performance which is higher than 50% in 7 out 10 training/test sets. In 2 out of the 10 training/test sets we obtained a performance higher than 82%. Moreover, our Miss Universe automatic prediction system was able to recognise the winner for the evening gown competition for years 1998 and 1999, which explains the higher performance for those years as in NDCG top ranked instances are considered more important. The predicted winner is also found in the top 3 for 5 out 10 versions of Miss Universe (2010, 2007, 1999, 1998, and 1996).

VII. MAIN FINDINGS AND FUTURE WORK

In this work, we have present a promising approach to automatically detect the winner during the evening gown competition of Miss Universe. To this end, we have created

TABLE II: NDCG for each year using best settings for SFV-PCA

Year	2010	2007	2003	2002	2001	2000	1999	1998	1997	1996
NDCG	77.71 %	78.95 %	52.97 %	62.78 %	44.37 %	44.71 %	82.93 %	87.02 %	51.69 %	77.38 %

a new dataset comprising 10 years of the evening gown competition selected from 1996 to 2010. We addressed this problem using action analysis techniques. We defined two problems that are of potential interest for the beauty pageant industry and the fashion industry. In the former problem, we are interested in predicting the winner of the competition, which can be also of interest for specialised betting sites. The fashion industry can have an innovative automatic system to compare two catwalks that can be used as training system for amateur models. Our system for predicting the winner of the evening gown competition shows we are able to rank the winner in the top 3 best predicted scores in 50% of the cases.

For future work, we propose to enlarge the dataset, to extend to the swimsuit catwalk competition, and the use of pose features. The current dataset can be enlarged using other Miss Universe versions, other beauty pageant competitions, and catwalks from international fashion trade shows. Given that scores are not always publicly available, an online competitive Catwalk rating game can be designed similar to the style rating game called *Hipster Wars* [23]. With this online game it would be possible to crowd source reliable human judgements of catwalks. The swimsuit catwalk competition together with the evening gown competition are critical to the selection of the next Miss Universe. For the swimming competition, other attributes apart from the catwalk would be needed to take into consideration such as good muscle tone, body proportion, body fat, body shape, and fitness. All those attributes are also visual attributes. Pose is an important attribute for catwalks. We envisage that pose-based Convolutional Neural Network features in conjunction with IDT can increase our system performance. This combination has been recently shown to be effective for action recognition [14].

Finally, we note that this work can be extended to other applications that require action assessment. For instance, patient rehabilitation and high performance sports. In both cases, an automatic system able to evaluate the progress of a patient or an athlete would be valuable.

REFERENCES

- [1] List of beauty pageants. http://en.wikipedia.org/wiki/List_of_beauty_pageants.
- [2] Miss Universe. <http://www.missuniverse.com/>.
- [3] Miss Universe in Wikipedia. http://en.wikipedia.org/wiki/Miss_Universe.
- [4] Polished by Donna. <http://www.polishedbydonna.com/>.
- [5] H. A. Abdul-Azim and E. E. Hemayed. Human action recognition using trajectory-based representation. *Egyptian Informatics Journal*, 16(2):187–198, 2015.
- [6] S. Ali and M. Shah. Human action recognition in videos using kinematic features and multiple instance learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(2):288–303, 2010.
- [7] X. Amatriain, A. Jaimés, N. Oliver, and J. M. Pujol. *Recommender Systems Handbook*, chapter Data Mining Methods for Recommender Systems, pages 39–71. 2011.
- [8] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [9] Z. Cao, T. Qin, T.-Y. Liu, M.-F. Tsai, and H. Li. Learning to rank: From pairwise approach to listwise approach. In *International Conference on Machine Learning*, pages 129–136, 2007.
- [10] J. Carvajal, C. McCool, B. C. Lovell, and C. Sanderson. Joint recognition and segmentation of actions via probabilistic integration of spatio-temporal Fisher vectors. In *Lecture Notes in Computer Science (LNCS)*, Vol. 9794, pages 115–127, 2016.
- [11] J. Carvajal, C. Sanderson, C. McCool, and B. C. Lovell. Multi-action recognition via stochastic modelling of optical flow and gradients. In *Workshop on Machine Learning for Sensory Data Analysis (MLSDA)*, pages 19–24, 2014.
- [12] O. Chapelle and S. S. Keerthi. Efficient algorithms for ranking with SVMs. *Information Retrieval*, 13(3):201–215, 2009.
- [13] W. Chen, T. yan Liu, Y. Lan, Z. ming Ma, and H. Li. Ranking measures and loss functions in learning to rank. In *Advances in Neural Information Processing Systems 22*, pages 315–323, 2009.
- [14] G. Chéron, I. Laptev, and C. Schmid. P-CNN: Pose-based CNN features for action recognition. In *International Conference on Computer Vision (ICCV)*, pages 3218–3226, 2015.
- [15] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. Liblinear: A library for large linear classification. *Journal of Machine Learning Research*, pages 1871–1874, 2008.
- [16] F. Gholami, D. A. Trojan, J. Kövecses, W. M. Haddad, and B. Gholami. Gait assessment for multiple sclerosis patients using Microsoft Kinect. arXiv preprint 1508.02405, 2015.
- [17] O. K. Gross, Y. Gurovich, T. Hassner, and L. Wolf. Motion interchange patterns for action recognition in unconstrained videos. In *European Conference on Computer Vision (ECCV)*, 2012.
- [18] K. Guo, P. Ishwar, and J. Konrad. Action recognition from video using feature covariance matrices. *IEEE Transactions on Image Processing*, 22(6):2479–2494, 2013.
- [19] Z. Hao, Q. Zhang, E. Ezquierdo, and N. Sang. Human action recognition by fast dense trajectories. In *ACM International Conference on Multimedia*, pages 377–380, 2013.
- [20] T. Joachims. Optimizing search engines using clickthrough data. In *International Conference on Knowledge Discovery and Data Mining*, pages 133–142, 2002.
- [21] T. Joachims. Training linear SVMs in linear time. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 217–226, 2006.
- [22] H. Kataoka, K. Hashimoto, K. Iwata, Y. Satoh, N. Navab, S. Ilic, and Y. Aoki. Extended co-occurrence hog with dense trajectories for fine-grained activity recognition. In *Asian Conference on Computer Vision (ACCV)*, pages 336–349, 2014.
- [23] M. H. Kiapour, K. Yamaguchi, A. C. Berg, and T. L. Berg. Hipster wars: Discovering elements of fashion styles. In *European Conference on Computer Vision (ECCV)*, pages 472–488, 2014.
- [24] C. P. Lee and C. J. Lin. Large-scale linear RankSVM. *Neural Computation*, 26(4):781–817, 2014.
- [25] H. A. Nguyen and J. Meunier. Gait analysis from video: Camcorders vs. kinect. In *International Conference on Image Analysis and Recognition (ICIAR)*, pages 66–73, 2014.
- [26] X. Peng, C. Zou, Y. Qiao, and Q. Peng. Action recognition with stacked Fisher vectors. In *European Conference on Computer Vision (ECCV)*, pages 581–595, 2014.
- [27] F. Perronnin, J. Sánchez, and T. Mensink. Improving the Fisher kernel for large-scale image classification. In *European Conference on Computer Vision (ECCV)*, pages 143–156, 2010.
- [28] H. Pirsiavash, C. Vondrick, and A. Torralba. Assessing the quality of actions. In *European Conference on Computer Vision (ECCV)*, pages 556–571, 2014.
- [29] L. Pishchulin, M. Andriluka, and B. Schiele. Fine-grained activity recognition with holistic and pose based features. In *German Conference on Pattern Recognition (GCPR)*, pages 678–689, 2014.
- [30] T. Qin, T.-Y. Liu, J. Xu, and H. Li. Letor: A benchmark collection for research on learning to rank for information retrieval. *Information Retrieval*, 13(4):346–374, 2010.
- [31] M. Rohrbach, S. Amin, M. Andriluka, and B. Schiele. A database for fine grained activity detection of cooking activities. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1194–1201, 2012.
- [32] J. Sánchez, F. Perronnin, T. Mensink, and J. Verbeek. Image classification with the Fisher vector: Theory and practice. *International Journal of Computer Vision*, 105(3):222–245, 2013.
- [33] C. Sanderson and R. Curtin. Armadillo: a template-based C++ library for linear algebra. *Journal of Open Source Software*, 1:26, 2016.
- [34] F. Wang, E. Stone, M. Skubic, J. M. Keller, C. Abbott, and M. Rantz. Toward a passive low-cost in-home gait assessment system for older adults. *IEEE Journal of Biomedical and Health Informatics*, 17(2):346–355, 2013.
- [35] H. Wang and C. Schmid. Action recognition with improved trajectories. In *International Conference on Computer Vision (ICCV)*, 2013.