Objective Sexiness? Using the Statistics of Feature Extraction Algorithm Outcomes to Determine the Sexiness of Images of Nude Women during Pornography Casting

Jakub Binter^{1,2} * Daniel Řiha¹, and Hermann Prossinger³

¹ Dept. of Anthropology, Charles University Prague, Czech Republic

² Dept. of Philosophy and History of Science, Charles University in Prague, Czech Republic

³ Department of Evolutionary Anthropology, University of Vienna, Vienna, Austria

*E-Mail: Jakub.binter@fhs.cuni.cz

Abstract. In this paper, we pursue the goal of finding an objective measure of sexiness of the naked female in an (adequately standardized) image of her while she is intent on eliciting sexual desirability (arousal) in the presumed viewer; we compare it with the objective measure of sexual desirability of her face when posing clothed. The data set consists of images of 15 women posing for selection in a casting (selection) session for participation in pornographic films; we can thus confidently assume that each female wants to appear as sexy as possible. We use AI-based feature extraction algorithms, implemented in software, to analyze feature vectors of the images. The women's apparent sexiness can be determined via statistical analysis. We argue that the feature extraction software we use supplies a clustering of these image feature vectors, implying an objective measure of the women's sexiness. We use such clustering, which implies a software-derived rating, but not ranking, when contrasting the ratings of these images by 50 men and 50 women. We argue that, as with the case of a protractor used to measure angles, objectivity can be derived by feature extraction, but not ranking by various human raters. Rather, the statistics of ratings by the men and women inform us of their personal evaluation of sexiness. These statistics of objective sexiness can then be used to infer (statistical) properties of the raters' personal sexiness perceptions and arousal response against our proposed constructed reference standard.

Keywords: Image feature extraction, artificial intelligence, dimension reduction, clustering algorithms, sexual attractiveness, image analysis of nudes, pornography film casting, sexual arousal.

1 Introduction

A large fraction of image analysis software is designed to follow as closely as possible the seeing mechanisms and visual perception processes by humans. There is, however, a caveat: the human visual system has flaws. Consider three scenarios. The first deals with the poetic description *Phaenomena* (written between 276 and 240 B.C.) of the constellations by Aratus of Soli. Based on previous work by Ancient Greek astronomers, he described clusters of stars as constellations, named them and ascribed properties to them in poetic forms; most notably those of the signs of the Zodiac, used profusely by astrologers since then. The second deals with asking a volunteer to use a ruler to draw two straight lines that are as perpendicular to each other as possible in his/her assessment. The third deals with the identification, by a radiologist, of the presence or absence of a malignant tumor in an x-ray of a female human mammary gland. All three cases deal with image features and their ‘subjective’ identification by the human visual system.

In the first case, it is easy to dismiss possible objectivity. While there is a long artistic tradition about the figurative appearance (the Greek word *Phaenomena* literally means appearances) of each constellation, and art historians have developed chronologies of how the constellation figures change over time [1], we doubt that serious thinkers would argue that these or other apparent clusters of stellar positions on the celestial sphere objectively show images. In fact, with the onset of publishing star atlases in the 16th Century, the debate where the lines separating constellations occur occasionally became quite heated, because the lines of demarcation were circles of celestial longitude and celestial latitude and could not be the boundaries of the constellation image as conceived and perceived by Aratus and subsequent astrologers.

The second case allows a straightforward reduction to objectivity. After drawing the two straight lines perceived by the person to be perpendicular, one uses a protractor to determine how close to 90° the angle between the intersecting lines is. We suspect no one will argue about the objectivity of the protractor. Whether the protractor reliably, reproducibly, and accurately measures an angle’s magnitude is a technical debate involving measurement standards and quality control institutions of the country in which the protractor was manufactured, not whether the person drawing the two lines is more reliable in his/her assessment of perpendicularity than the protractor is.

The third case is now often celebrated by the AI community as a success story of using image processing software in medical diagnostics. The use of AI image analysis software as a diagnostic tool was initially dismissed as less reliable than a radiologist’s trained eye (by the radiologists, we note). Meanwhile, the situation has reversed dramatically: current statistical findings show that AI image analysis is far superior [2] to diagnosis based on human vision (because there are less false negatives).

The ‘moral’ derived from the above three cases is: while it may oftentimes be useful to design image analysis software that attempts to mimic human vision as closely as possible, it is also often imperative to design software that achieves a degree of objectivity which could be at odds with the human vision system. Clearly, the claimed objectivity of the software must then be determined by means other than by referring to

what a human ‘sees’. AI deals with this objectivity issue by applying statistical techniques to training/test set outcomes, a method we use here.

A fourth scenario tells us something about the evolution of humans and their partnership finding strategies prior to copulating. Consider: a young man and young woman are watching another woman of similar age for the first time. She — unavoidably and unintentionally — makes an impression on both. If she were (physically and/or sexually) attractive to the young man, she could become a date for him and also a potential rival to his present female partner. The central field of view of human vision orients itself towards biologically relevant markers of phenotypical quality of an individual someone communicates with [3]. Dominance, formidability, attractiveness, and levels of steroid hormones have been shown to be relatively correctly assessed by raters [4][5]. Thus, when assessing specific properties of another human, possibly a mate or a rival, the facial cues are evaluated for what they signal, as signaling theory has evidenced [6][7]. The facial features/structures imply variations in the above-mentioned characteristics. The viewer relies on the entire face being evaluated holistically. The existence of this bias can be demonstrated by what happens after turning the presented facial photographs upside down; then raters drastically fail in their ability to deduce information [8].

An alternative approach can be used to establish features that play roles in such recognition and assessment processes.

We propose using image analysis software to determine the sexiness of human females via images of them in the nude, as well as of their faces when they are clothed. In a later publication, we then will compare these sexiness measures (described below) with the sexiness and physical attractiveness ratings male and female human assign to these images. Much like the scenario of using a protractor to determine how well a person succeeded in estimating a 90° angle, we argue that we can determine how ‘well’ a human can assess the sexiness of the naked female appearance by comparing the human rating with objectively determined sexiness that is outputted by image feature extraction software. We add that we are cognizant of, and want to enter the debate about, whether ‘sexiness’ can even be considered objective. Our position in this paper is: ‘yes’ it can be, and the inability of humans (in particular the males) to be objective in sexiness assessment is a driver of human evolution. If there were hardly any error with rating sexiness as a degree of arousal as one (very important) variable (of several) when copulating with a partner in a multivariate statistical choosing scenario, then arousal variability is lost; the pool of sexy partners to choose from would drastically decrease the response flexibility of reproductive partners in a species. Evolution would not function optimally, because then only random fluctuations would be responsible for partner choice.

For the pornography film industry, this lack of ability by men (primarily, but also by women) to objectively assess the sexiness of naked women (among other criteria) is literally vital; the industry must offer a wide range of female appearances alluding to their sexiness so that the potential customer’s desire for sexual arousal is heightened prior to consuming this industry’s products. There is no single female nude that arouses all consumers equally strongly (a well-known observation [9] that is also supported by

analyses of our data set). In fact, one reason for the casting mechanism in the pornographic film industry is based exactly on this observation: to avoid choosing only those females as actors in pornographic films that are perceived as sexually arousing for the director and/or the producer, but not necessarily for everyone in the target audience, as the arousal spectrum is not uniform. In pornographic videos, the females' sexiness must be (literally) in the eye of the beholder/consumer.

2 Materials

2.1 Photographs

15 women participated in a casting session hoping to be chosen for acting in pornography films/videos. Prior to casting while in action, they posed for two photographic stills: naked (except occasionally with minor jewelry) and solely the face while clothed. The females posed nude in front of a white, mildly textured wall, standing on a dark brown wooden floor, in full frontal view, feet close together, arms along their sides, and illuminated by standardized lighting that was the same for all participants. A photograph of each casting participant's face, while she was clothed, was also taken in the same room under the same lighting. The face images were not electronically extracted from the nude images.

3 Methods

3.1 Collection of Ratings

A counterbalanced two-block design was employed in which participants rated the face when the participant was clothed, and the nude (body together with face). The resulting stimuli were entered on two separate Likert-type scales: (1) physical attractiveness (from 1=A not physically attractive, to G=7 highly physically attractive), and (2) sexual desirability (from 1=A not sexually desirable, to G=7 highly sexually desirable). An option to skip a rating was also offered in the unlikely event that any of the female stimuli were recognized by a rater. Definitions of the two scales were presented to participants on a display screen at the start of the procedure; the descriptor "Physical attractiveness" is "to what extent this stimulus conforms to the notion of ideal physical appearance", and "Sexual desirability" is "to what extent this stimulus could evoke feelings of arousal, irrespective of physical appearance."

3.2 Bayesian Likelihood

To support the claim that the male raters distinguish between sexiness and physical attractiveness, we use the machinery of Bayesian statistics [10][11]. If s is the probability ($0 \leq s \leq 1$) that sexiness of a female is rated higher than her physical attractiveness, then $evidence = \dots s^{n_1}(1-s)^{n_2}$, where n_1 is the number of occurrences that sexiness is rated greater than physical attractiveness, n_2 is the number of occurrences

that physical attractiveness is rated greater than the female's sexiness, and $\dots$ is a constant to be calculated by integration. The likelihood function $\mathcal{L}$ is the probability density function of the probability s , which is a random variable in Bayesian statistics. One consequence of Bayes Theorem is that $\mathcal{L}_{\text{post}} = \text{evidence} \times \mathcal{L}_{\text{prior}}$. In this paper, we assume a Bayesian prior $\mathcal{L}_{\text{prior}} = \dots s^0(1-s)^0 = 1$, because we have prior evidence that women have both a physical attractiveness and a sexiness that can be rated; an observation based on the sexuality-driven copulation behavior of *Homo sapiens* evidenced during a more than $\frac{1}{4}$ of a million years' time span. The posterior likelihood $\mathcal{L}_{\text{post}}$ is a Beta function

$$\mathcal{L}_{\text{post}} = \frac{\Gamma(n_1+n_2+2)}{\Gamma(n_1+1)\Gamma(n_2+1)} s^{n_1}(1-s)^{n_2} \quad (1)$$

where $\Gamma(\dots)$ is the Gamma function [12]. The ML (most likely) probability $s_{\text{ML}} = \frac{n_1}{n_1+n_2}$ is the mode of the likelihood function and the $\text{HDI}_{95\%}$ is the interval $s_1^{s_2}$ about the mode so that $\int_{s_1}^{s_2} \mathcal{L}_{\text{post}}(u) du = 0.95$ that gives the 95% probability that the 'true' ML probability lies within this interval (a more elaborate derivation is given in [11]). The probability that the obtained result is due to chance can be estimated by $\int_0^{\frac{1}{2}} \mathcal{L}_{\text{post}}(u) du$ or $\int_{\frac{1}{2}}^1 \mathcal{L}_{\text{post}}(u) du$, depending on whether $\text{mode} > \frac{1}{2}$ or $\text{mode} < \frac{1}{2}$.

3.3 Feature Extraction

Modern feature extraction applied to images uses unsupervised machine learning algorithms, usually neural networks (as we do in this investigation). We supply as input to these algorithms 15 images of (a) full nude images in frontal posture, and, separately, (b) almost frontal images of the participants' faces when they were clothed (parts of clothing are visible in these face images; the major portion of the image is the face, however). The learning algorithms are trained (via repeatedly randomly partitioning the data into training and test sets) many times on each image. We thus obtain a feature vector for each image in each of the two input sets: i.e. 15 feature vectors for the nude images and 15 feature vectors for the face images.

3.4 Further dimension. reduction and clustering

We use a further dimension reduction step, after the feature extraction algorithm has supplied (in our investigation here) each image with a 14-dimensional (14D) feature vector. Each set of 15 images results in a 15×14 matrix, one matrix for the nude images and one for the faces. (a) We use an extraction algorithm that searches for the best linear combination of the 14D feature vectors to extract a 2D vector of sexiness for each female. The resulting 2D feature vectors inform us of the diversity of sexiness of the females (more specifically: of their images while posing).

We apply a clustering algorithm to this dimension-reduced set of vectors to find a suitable number of clusters (one or two clusters do not make sense; any more than 14 or 15 clusters do). For practical, empirical purposes, we chose seven clusters; then we

have a straightforward sexiness clustering comparable to the number of categorical rating options for the 50 rating men (their rating scale is A–G; corresponding to 1–7).

3.5 Predicting sexiness

We extract a feature vector of a further nude model that was not one of the 15; reduce the dimensions of this feature vector to 2D and locate its position in the clustering graph. The nearness of this 2D feature vector to the convex hull of a cluster or its belonging to the cluster is a **measure** of the objective sexiness of the nude image.

Fig. 1. Frequencies and Bayesian likelihood $\mathcal{L}_{\text{posterior}}$ that Female No. 12 is sexier than physically attractive, according to the ratings by 50 males. **(a)** The bar chart of frequencies in the sexiness-physical attractiveness rating grid. This bar chart shows that few frequencies are on the diagonal plane where sexiness and physical attractiveness are equal. The frequencies when males rate Female No. 12 as sexier than physically attractive occur more often (more and higher bars on the “sexiness” side of the diagonal plane). **(b)** The posterior Bayesian probability that a male will rate Female No. 12 sexier than physically attractive. HDI_{95%} is the “highest density interval with a confidence of 95%” [11]; this interval is not symmetric about the ML probability, because the Beta distribution is asymmetric. The mode (of the posterior likelihood) is 75.9%, within an interval of 58.9%–88.6%, clearly indicating that there is a high probability that (future) males will rate this female as sexier than physically attractive. The probability that this result is due to chance is 0.3%

4 Results

4.1 Distinction between sexiness and physical attractiveness

Fig. 1 shows that raters overall do distinguish between physical and sexual attractiveness, with a high reliability. Table 1 shows that very few females posing nude are more likely to be physical attractive than sexually attractive.

Table 1. Bayesian probabilities of sexiness of the 15 females (nude) rated by the 50 males. The second column from the left lists how many males rated a female with a sexiness rating higher than a physical attractiveness rating (in contrast to a physical attractiveness rating higher than a sexiness rating; ties were not included in this Bayesian analysis). The $HDI_{95\%}$ is the highest density 95% interval, as explained in the Methods section. The rightmost column lists the probability that the observed sexiness or non-sexiness is due to chance, as explained in the Methods section. We observe that males rate 10 females (i.e. $\frac{2}{3}$) as sexier than attractive, if we consider $P_{\text{chance}} < \sim 5\%$ (we do not consider significance levels as reliable; see [13]). There is a marked gap: only two P_{chance} are above 10%; these two females have been rated predominantly physically attractive. For 11 females $HDI_{95\%}$ does not overlap $s = 0.5$.

Female No.	Frequency Sexy > Attractive	$HDI_{95\%}$			P_{chance} (%)
		$s_1(\%)$	$s_{ML}(\%)$	$s_2(\%)$	
1	20	55	73	87	1.0
2	8	13	28	47	1.4
3	20	48	66	81	4.9
4	17	47	67	83	5.4
5	20	58	76	89	0.5
6	24	55	72	85	0.7
7	8	12	25	42	0.4
8	12	24	41	59	17.0
9	18	71	89	98	0.02
10	18	66	85	96	0.07
11	21	46	62	78	8.1
12	23	59	76	89	0.3
13	12	54	79	94	1.8
14	13	37	57	76	26.0
15	25	59	75	88	0.2

4.2 Frequencies of sexiness ratings of females

Fig. 2. The sexiness (sexual desirability) of the females who participated in a porno casting, as rated by 50 male raters on a rating scale of A–G (1–7). **(a)** Five nude females (Nos. 4, 5, 9, 10, and 13) have very low ratings. No female participant has exclusively high ratings. Furthermore, in contrast to Female No. 4 (which was very rarely rated above E and never G) a few males find Females No. 5, 10 and 13 highly sexually desirable, despite the numerous low ratings by the other males. Because these females were given high ratings by close to 10 % of the male raters, we reject the argument that the raters made errors when entering their ratings into the data base. We infer, rather, that the sexiness rating of a female by a male rater is not an objective quantity; obviously, some men find some females sexually very desirable, even though a large fraction of many men strongly disagree with such high ratings. **(b)** Four female faces were perceived as being particularly low in sexiness (‘unsexy’): Females Nos. 5, 9, 10, and 13 (compare with Fig. 2a). Further evidence for variability of rated sexiness of the female face (by males) is evidenced in Female Nos. 9, 10, and 13 (some males rate them at F and No. 9 even up to G) and in Females Nos. 1 and 15 the distribution of ratings is close to uniform. The sexiness of Female No. 4’s full nude image (Fig. 2a) contrasts with the sexiness of her face (almost every male rates her nude image A; the mode of her face rating is C, and A occurs rarely).

Fig. 1a demonstrates the frequencies of ratings assigned by the 50 male raters to the image of Female No. 12 when nude. Sexiness is rated more frequently with higher ratings than physical attractiveness (Fig. 1a). The resulting likelihood function (Fig. 1b) supports the conclusion that, for all the 50 male raters combined, she is sexier than physically attractive with a ML probability 76%. The $HDI_{95\%}$ uncertainty interval for this mode is listed in Table 1. We do not use this result further in this paper, as we attempt to find an ‘objective’ measure of sexiness, which we do not derive from the ratings by the males. Fig. 1 clarifies how one could proceed to analyze a population statistic of the ratings of any nude female. In a sense, Fig. 1 shows how we cannot

proceed to find sexiness objectively. It also demonstrates that the physical attractiveness of a female in a nude image is to be distinguished from its sexiness.

In order to support our argument that sexiness must be considered different from physical attractiveness, irrespective of whether it can be objectively defined, we analyze the frequencies of ratings by the 50 male raters for each female. $\frac{2}{3}$ of the females are rated as sexier than attractive. Table 1 lists further analyses outcomes, based on Bayesian statistics. Because we do not analyze the sexiness of the females as rated by the males, we only present these results here to demonstrate that, for the 50 male raters, the females did indeed pose in such a way as to provoke a high sexiness rating by some — indeed, many — raters. To some extent, we can consider the results demonstrated in Table 1 as evidence for how well the female casting participants succeeded — in a population statistics sense.

4.3 Clusters

Fig. 3. The sexiness clusters of (a) nude female images and (b) face images when clothed as determined by feature extractions of the images and dimension reduction. Each cluster is color-coded; the convex hulls of each cluster are a guidance for the coordinates of the feature vectors of the sexiness of females in the clusters. The number identifying the female is from the list of images; there is no information encoded in this number, nor in the color coding of the cluster convex hulls. (a) Of the seven clusters, only one is a singleton, and three contain two females. Female No. 16 is that of a further female whose sexiness is to be determined by the sexiness of the rated 15 females; further details are elaborated in the text. (b) The clustering of the faces differs from those in sexiness of the full nude (Fig. 2a). One cluster (Female Nos. 1, 3, 12, 14) contains four females with comparable sexiness. The sexiness of Female No. 15's face is markedly different from that of all the other females' faces.

For the nude images, none of the seven clusters contains more than 3 female images (Fig. 3). We find that the elements of the seven clusters demonstrate that dimension reduction to 1D is not warranted. That sexiness is not a scale on a one-dimensional

linear manifold is, in our opinion, a major contribution to attempting to determine sexiness objectively: objective sexiness cannot be found by looking for some linear manifold on which to map the sexiness scores.

4.4 Feature (image) distances

In technical parlance, feature distance (image distance) is the Euclidean norm between two feature vectors. Feature distances do, of course, change after the implementation of feature reduction algorithms: that is the point of doing so. With the use of clustering, we can determine how close the sexiness of a female in the sample is to the sexiness of another female in the sample by calculating the feature distance.

For the nude images, we note (Fig. 3) that the sexiness feature vectors of females in some clusters are almost identical (Females Nos. 5 and 15, for nudes; Female Nos. 5 and 10 for faces). In all, we infer that classifying sexiness into seven classes seems heuristically meaningful; each sexiness class, determined by the convex hull of the clustered sexiness, appears reasonable. No cluster contains a large fraction of the 15 feature vectors, and only one female (No. 13) is a singleton in the nude sexiness (Fig. 3a).

We observe a comparable finding for the images of the females' faces. The sexiness of Female No. 15 is very different (and very far from) those of the others. The largest cluster contains the sexiness feature vectors of four females, and there are two singletons. Three clusters consist of pairs, and for two of these the feature distances from each other within the cluster are small.

4.5 Prediction

In Fig. 3, we show where a further female's nude image (Female No. 16) is located in the 2D sexiness graph. This female's (nude) sexiness clusters with the sexiness of Females Nos. 4 and 11.

5 Discussion

5.1 Relevance of statistical methods

Considerable research has been invested in finding algorithms that are suitable for feature extraction from images. Mlsna & Sirakov [14] presented a history, an overview, and a review of some important features that are to be extracted from an image: segmentation, convex hull models, shape features, etc.; all these for each color channel. Since then, progress has been breathtakingly rapid, due to the now commonplace implementation of the AI paradigm of repeatedly randomly splitting the data set into a training and test set and the advances in designing neural net algorithms as learning tools. The data sets that are used as training/test sets and also used to fine-tune the commercially available algorithms have also grown enormously in the last few years: the software package MATHEMATICA® (which we use) has more than ten million

images (only a small fraction — about 60 thousand — are face images) and thus can be considered a trustworthy base for relying on the algorithms’ applicability.

5.2 Objective Sexiness

Calculating the Euclidean norm of the sexiness feature vector from the origin and defining it as the apparent sexiness of the female, both for the face and for the nude image, is problematic and fraught with logical inconsistencies. If this norm were a reliable sexiness measure, then a feature reduction to 2D feature vectors should not be implemented, but rather a 1D feature reduction instead. Then we could use the Euclidean norm and assign sexiness on a linear scale. If the centroids of the clusters were nearly on a straight line (not necessarily parallel to one of the coordinate axes) or some smooth 1D manifold embedded in 2D, then a 1D feature reduction would not only be warranted, but obligatory. One insight we gain from implementing the methodology we explicate and the outcomes we present is the necessity of not assessing sexiness on a one-dimensional manifold. For this reason, the clustering in Fig. 3 shows comparable sexiness (note, however, the different scales on the axes); the clusters themselves should not and cannot be ranked. We found the clusters by *K*-means clustering [10][15]. *K*-means clustering is an algorithm that finds elements that are closer together within a cluster than between elements of other clusters. This may seem to be a ranking algorithm; it is not, because the clusters are not collinear in any sense; in particular, their centroids are not — not even on a smooth manifold (Fig. 3).

The sexiness of an image of a further female (i.e. one not used in the clustering derived from the 15 feature vectors) is very close to Female No. 4’s sexiness (Fig. 3).

We find that sexiness of faces clusters differently from sexiness of images of the nude body. Perhaps unexpected, the variability of sexiness of the fully nude females’ images is smaller than the variability of the sexiness of their faces when clothed. In this presentation, we do not segment the full nude image into sections (for example: face, then from the mean of the highest shoulder points to navel as one horizontal rectangle and then navel to knees as another) so as to determine how such sections vary and cluster in sexiness. We consider this approach too fragmentary, because the data set consists of only 15 female participants and we cannot, in the next part of the investigation, compare the objective sexiness (the 2D feature vectors) with the ratings by the males and females. It is, of course, possible to segment the nude images to determine sexiness feature vectors of each segment, but it would not supply much insight in this context without being able to refer to how humans rate sexiness. (A common adage in the AI community: “Use AI when your understanding is poor and the data sets are rich.”)

We repeat the argument why we consider the feature vector outcomes (in 14D) as measures of objective sexiness, rather than of attractiveness: the female participants participated in a casting for pornography clips/movies/videos; it was their intention to appear as arousing and sexually desirable as possible (in their self-perception of what constitutes ‘arousing’) to the prospective male viewers of the clips they desire to act in. Their presentation of their sexiness can, of course, misfire: they might have appeared sexier to the male viewers if they had presented themselves differently. One attempt at

partially controlling for this conundrum is the requirement imposed by the film crew involved in the casting; the females were to stand upright in frontal posture — the shoulder-to-shoulder axis not oblique to the pelvis, for instance, and with legs closed or nearly closed (and parallel to each other and to the wall). The casting procedure is not complete with the making of these images (‘stills’); the prospective participants had to act in short pornographic clips, invariably involving copulations, etc. with a male actor. Nonetheless, determining the sexiness of the images is still very valuable. Potential clients of a pornographic video on a DVD, say, first look at the cover; there, the sexiness of the ‘stills’ is important (“A client buys the porno DVD by its cover.”). The images of the faces we used in this analysis were not segmented from the nude images. In fact, some females had their head slightly askance when they posed for the facial photographs; most of those with long hair draped them over one shoulder (rather than letting it hang down their back), creating a further asymmetry. The feature extraction algorithm is little driven by such aspects of imagery, we found; the males’ ratings were probably not either. We ascribe the differences between the clusters ‘nude’ and ‘face’ to the females’ desires to appear sexy in the nude stills and beaconing sexual desirability in facial expression when clothed.

A seemingly obvious application of the investigation of objective sexiness is the programming and implementation of recommendation software: letting a potential client select pornographic clips/videos/movies based on the preferences spectrum of this prospective male (and sometimes: female) client. Writing such application software is only possible after knowing the preferences, as we show in Fig. 2. However, the analyses we present here are a prerequisite; the objective sexiness is the independent predictor variable in a (to-be-programmed) mapping function that maps a chosen/preferred sexiness onto the client’s desires.

One potential application may also loom after completion of an AI investigation of sexiness of nude images: the creation of avatars in an illusionary (or fantasy-type) pornographic film/video. Products using nude avatars are already on the market, but we claim that using the analyses of objective sexiness can enhance the effectiveness of avatar construction. We have demonstrated how to implement avatar construction that focusses on sexiness. The nude image of Female No. 16, as a surrogate for an avatar, is close to the feature vectors that have already been clustered. Again, because the sexiness of a female participant/avatar is not the same for every potential customer and subsequent consumer, several avatars must be created with a spectrum of sexiness. An assortment of avatars is obligatory; their image feature vectors should lie within or close to the convex hulls (Fig. 3). We have demonstrated statistical properties of the sexiness spectrum of females and to-be-constructed avatars here.

6 Conclusion

We claim that we can use image feature extraction algorithms, incorporating the AI paradigm of test and training sets to produce an objective sexiness score for an image of a female posing for participation in a pornographic clip/film/video. The assumption we make is that the female desires to appear as sexually desirable and arousing to the

male viewer while the photograph is taken. We first use a feature extraction and then a feature vector reduction algorithm to generate a 2D scatter plot of the sexiness scores, followed by a clustering algorithm to generate clusters of these objective sexiness scores. Female sexiness scores within each cluster are comparable and every score within a cluster is close (in the feature vector distance sense) to that of her competitors in the casting. We found that a sexiness score is two-dimensional and cannot be projected onto a one-dimensional smooth manifold.

The raters, however, were asked to rate their perceptions of sexiness of each female's nude image and face image on a linear scale from A–G (1–7) (Fig. 2). Asking the raters to do so, rather than rate each image in a 2D plot is reasonable, because humans are used to ranking on linear scales in 1D. We infer, therefore, that objective measures of sexiness should not be designed so as to mimic human ratings of sexiness.

References

1. Boll, F. (1903) *Sphaera: Neue Griechische Texte und Untersuchungen zur Geschichte der Sternbilder*. Teubner, Leipzig.
2. McKinney, S.M., ... (31 additional authors) ..., Shetty, S. International evaluation of an AI system for breast cancer screening. *Nature* 577, 89–94 (2020). <https://www.nature.com/articles/s41586-019-1799-6>
3. Wells, T.J., Dunn, A.K., Sergeant, M.J.T., Davies, M.N.O. Multiple signals in human mate selection: A review and framework for integrating facial and vocal signals. *Journal of Evolutionary Psychology*, 7(2), 111–139 (2009).
4. Thornhill, R., Grammer, K. The body and face of woman: One ornament that signals quality? *Evolution and Human Behavior* 20(2), 105–120 (1999).
5. Weeden, J., Sabini, J. Physical attractiveness and health in Western societies: A review. *Psychological bulletin* 131(5), 635 (2005).
6. Scott-Phillips, T. C. (2008). On the correct application of animal signaling theory to human communication. In: Smith, A.D.M. et al. (eds.), *The Evolution of Language: Proceedings of the 7th International Conference on the Evolution of Language* (pp.275–282). World Scientific, Singapore.
7. Scott-Phillips, T.C. Defining biological communication. *Journal of Evolutionary Biology*, 21(2), 387–395 (2008).
8. McKelvie, S.J. Emotional expression in upside-down faces: Evidence for configurational and componential processing. *British Journal of Social Psychology*, 34(3), 325–334 (1995).
9. Schaschek, S. *Pornography and Seriality: The Culture of Producing Pleasure*. Springer, New York, USA (2013).
10. Bishop, C.M. *Pattern Recognition and Machine Learning*. Springer Science+Business Media, New York, USA (2006).
11. Kruschke, J.K. *Doing Bayesian Data Analysis. A Tutorial with R, JAGS, and STAN*. Elsevier, London, UK. (2015).
12. Davis, P.J. Chapter 6: *Gamma Function and Related Functions*. In: Abramowitz, M., Stegun, I.A. (eds.) *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, New York, USA (1972).
13. Amrhein, V., Greenland, S., McShane, B. Retire statistical significance. *Nature* 567, 305–307 (2019).

14. Mlsna, P.A., Sirakov, N.M. Intelligent shape feature extraction and indexing for efficient content-based medical image retrieval. 6th IEEE Southwest Symposium on Image Analysis and Interpretation, (2004).
15. MacKay, D.J.C. *Information Theory, Inference, and Algorithms*. Cambridge University Press, Cambridge, UK. (2003).

Appendix: Software Details

We use MATHEMATICA[®] v12 from Wolfram Technology. The codes for the algorithms are proprietary and depend on internal tooling. However, the executions of the codes are available in this commercial software package. In particular, we executed three dedicated functions:

- (1) **FeatureExtract[image]** to extract the features of **image**.
- (2) **DimensionReduction[featureVector, 2, Method->"linear"]** to obtain 2D feature vectors for each female image ('nude' and 'face' separately), automatically choosing the best linear method ('best' being decided via AI learning methods).
- (3) **FindClusters[featureVectorSet2D, 7, DistanceFunction->EuclideanDistance]** to find 7 clusters using the *K*-means algorithm (the default in this command) and the Euclidean distance as **DistanceFunction** between the feature vectors in the two sets **featureVectorSet2D** (one set for the nude images and one set for the face images).

The graphics, the tallying, and the (Bayesian) statistical analysis have also been programmed (by the 3rd author) in MATHEMATICA[®].