

Document Version

Final published version

Licence

Dutch Copyright Act (Article 25fa)

Citation (APA)

van Zuijlen, M. J. P., Pont, S. C., & Wijntjes, M. W. A. (2020). Conventions and temporal differences in painted faces: A study of posture and color distribution. *IS and T International Symposium on Electronic Imaging Science and Technology*, 32, Article 14. <https://doi.org/10.2352/ISSN.2470-1173.2020.11.HVEI-267>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Conventions and temporal differences in painted faces: A study of posture and color distribution

Mitchell J.P. van Zuijlen, Sylvia C. Pont, Maarten W.A. Wijntjes

Perceptual Intelligence Lab, Faculty of Industrial Design Engineering, Delft University of Technology, Delft, the Netherlands

Abstract

The human face is a popular motif in art and depictions of faces can be found throughout history in nearly every culture. Artists have mastered the depiction of faces after employing careful experimentation using the relatively limited means of paints and oils. Many of the results of these experimentations are now available to the scientific domain due to the digitization of large art collections. In this paper we study the depiction of the face throughout history. We used an automated facial detection network to detect a set of 11,659 faces in 15,534 predominately western artworks, from 6 international, digitized art galleries. We analyzed the pose and color of these faces and related those to changes over time and gender differences. We find a number of previously known conventions, such as the convention of depicting the left cheek for females and vice versa for males, as well as unknown conventions, such as the convention of females to be depicted looking slightly down. Our set of faces will be released to the scientific community for further study.

Introduction

The human face is extremely informative; faces allow us to make visual inferences of an individual's identity, age, gender, health and emotional states. It is argued that face perception is the most developed skill in visual processing and has received much attention from a psychological viewpoint [1-3]. It is often neglected that artists have been interested in the perception – and depiction – of the human face for over a thousand years. These depictions of faces can be found in a large variety of cultures throughout history in a large variety of styles ranging from the Roman-Egyptian Fayum portraits to modern hyper-realistic portraits. Painters have become masters in depicting faces by employing careful experimentation using the relatively limited means of paints and oils. Many of these faces are now available to the scientific domain within recently digitized collections of various art institutions [4, 5]. Modern facial detection algorithms can be used to automatically detect and segment these faces from art collections.

We are interested in quantifying artistic conventions.

The depiction of faces is interesting within this context. For example, it has previously been shown that three-fourths views are much more common relative to a full-frontal view. This effect has been found in various media such as paintings, photographs, etchings and drawings [6] and even in contemporary visual culture [7,8]. Initially, it was suspected that this view might increase recognition accuracy [9], but this hypothesis has since then been rejected [10, 11]. McManus & Humprey (1973) found that painters are more likely to specifically depict the left cheek of the subject, especially for females, which has been interpreted to be the side of the face that is more emotionally expressive [13, 14]. Left cheek here refers to the left cheek of the subject being

painted (see figure 2). More of these conventions within paintings are known, for example, it was found that painters tend to horizontally center one eye within the canvas [15]. However, it is often unknown how conventions such as these change over time. In recent work, Carbon & Pastukov (2018) looked at the historical trend of one well-known convention: that light in art comes from the top-left. They found specifically that this convention starts with the early renaissance and that the convention stays relatively constant over time. Within this paper we would like to study such conventions, and see how and if these conventions change over time.

A second topic of interest here is human skin which is of interest to a variety of fields, such as art history [17], dermatology [18], optics [19, 20], perception [21-23], computer vision [24] and computer graphics [25, 26]. See [27] for a comprehensive overview of literature about the appearance of skin. Skin is one of the most depicted materials within art history. In order to study material depiction e.g. [28] the availability of skin patches is of substantial relevance.

In this paper we present an initial analysis on the pose metrics of the face, combined with some basic color statistics of the skin and further relate these statistics to gender and metadata about the painting's age. To summarize: the goal of the current paper is to study conventions within art related to the artistic depiction of the human face, specifically the pose and skin, and see if and how this change over time.

Table 1. Number of paintings per gallery collected.

Gallery	# of paintings
The J. Paul Getty	399
The Metropolitan Museum of Art	3222
The National Gallery of Art	2132
Museo Nacional del Prado	2032
Het Rijksmuseum	4672
NationalMuseum	3077

Methods

Paintings

We used images from six international art galleries: The Getty, The Metropolitan, The National Gallery of Art, The (Spanish) Museo Nacional Del Prado, the (Dutch) Rijksmuseum and the (Swedish) NationalMuseum. The collections for each of these galleries are available to the public on their respective websites. The links to these websites have been included in the references. We used web scraping, or API's when available to download each available painting in these collections. The majority of paintings we downloaded also possess metadata about the (estimated) year of production. In total we collected 15,534

paintings in this way. The number of paintings per gallery has been included in table 1. The paintings are predominately western art and as a result primarily depict Caucasian skin.

A number of paintings collected were monochrome. We found that a number of these were actually sketches or drawings, instead of paintings. Therefore, we decided to filter out monochrome images. We detected monochrome images by using a linear regression on the three-dimensional color data for the original images. After filtering these monochrome images, we had 14,873 paintings remaining.

Face detection and estimations

To detect faces within paintings we used the commercially available Face++ facial detection network, available on www.faceplusplus.com. With the localization of each face, the network marks the (x, y) coordinates of 83 landmarks, which are visualized in Fig. 1, that correspond to various facial features, such as the eyes, mouth, etc. The Face++ network detected a total of 11,659 faces.

We performed a visual inspection on three random subsamples of 100 detected faces each and found a 1% rate of false positives for face detection. We interpreted this as a negligible effect and applied no corrective measures.

The Face++ network gives estimates of the head pose, the age and the gender for each detected face. The exact algorithm Face++ used to generate these estimations are proprietary and are not known to the authors. First, the head pose is expressed in three angles of rotation: yaw (horizontal rotation around body midline), pitch (vertical rotations around ear-to-ear axis) and roll (rotations around the axes in the direction of the nose). Second, the age of each person behind the face is predicted in years. Visual inspection suggested that the age prediction was inaccurate and we decided not to use this metric in further analysis. Last, the network also made a prediction of gender for each detected face. Visual inspection of these data suggested it was accurate enough to include gender in further statistical analysis, but it should be noted that the network appears to perform below human levels.

Face++ was trained on a set of images, just like any other automated detection algorithm, which can introduce a number of detrimental effects. Any possible bias within this training set of images might be transferred to the results when testing another set of images, even if this second set of images does not contain this bias inherently. Furthermore, Face++ has been trained on natural images, that is, photographs. When interpreting results from networks such as Face++, these possible issues should be noted.

Color

The paintings were digitized by each gallery and it is likely that each gallery applied different color calibration and color management techniques. Furthermore, each gallery likely also made use of different cameras, which also has an influence on the color mappings. Regrettably, we do not have any color calibration data for these collections. The lack of such data could potentially invalidate results based on the color data, especially if analyzing individual color values. Therefore we analyzed a number of global color distribution metrics that can safely be assumed to be rather robust under such transformations. Moreover, we studied

primarily their development over time (i.e., as a function of the year of creation of the painting).




Figure 1. A face that was detected by the Face++ network, overlaid with the 83 landmarks defined by said network. The blue lines connect all landmarks labelled as face contour.

We assume that color calibration techniques are applied relatively consistently within different galleries, but not between galleries. If for example, one of the galleries applies a color calibration technique that captures all the painting images of that gallery with increased luminance relative to other galleries, we assume that this is consistent over their collection ranging over a certain period and thus over time. Here it should be noted that the periods covered by the different collections overlap. As such, gradual effects in luminance over time would therefore likely not be an effect of color management or calibration techniques.

Additional sources of deviations from the overall skin color are introduced by the background, as well as the eyes and mouth. Therefore, we filtered out the background using the Face++ landmarks designated as the face contour (the blue line in figure 1). Next, we removed the eyes and mouth by plotting black dots around the landmarks relating to the eyes and mouth, where each dot had a diameter set to the 10% of the width of the original bounding box containing the face landmarks. An example of an original face, as well as a filtered face can be seen in figure 2.

We transformed the remaining skin pixels into Lab color coordinates and then extracted the luminance values. Next, we calculated the mean and variance. Here the mean luminance would indicate the lightness of the skin while the variance captures the range of luminance, which is related to contrast. Furthermore, we wanted to capture the amount of colors used and tried to create a 'colorfulness' metric. We define our colorfulness metric by computing the surface area of all pixels in the ab-plane within a Lab space. We created a grid in this ab-plane, with a lattice constant of 0.02 and counted whether each bin was taken by at least 5 pixel values. We visualized an example in figure 3.



Figure 2. **Left:** A bounding box of the original face: a detail from "The Holy Family with Saints Francis and Anne and the Infant Saint John the Baptist" by Peter Paul Rubens (1902). **Middle:** The same face, but with the background removed using automated landmarks from the Face++ algorithm. **Right:** The same face, but this time also with the eyes and mouth features removed.

Analysis

For the analysis we used 8 characteristics: pose data (3), color data (3), year of creation of the painting, and the gender of the face, estimated by Face++. The pose data is expressed in pitch, roll and yaw angles and have been discussed above. For the color data we analyze the mean and variance of the luminance, as well as our own defined colorfulness metric.




Figure 3. The area occupied in the *ab*-plane of Lab space. This is defined as a measure of 'colorfulness', that is the amount of colors used by painters to depict skin.



Figure 4. **Top:** Pitch angles. **middle:** Roll angles. **Bottom:** Yaw angles. For each of the three angles of rotation the range from lowest on the left to highest on the right has been visualized




Figure 5. The angles of rotation illustrated. Pitch rotation would revolve around the green axis, where the face would look down or up. Roll rotation revolves around the yellow axis, which would result in tilting one's face to the left or right. Last, rotations along the blue axis lead to changes in yaw, for which the face looks towards the left or right.

Results and discussion

Orientation

First, we wondered in what direction the faces were looking. In figure 4 examples are shown of varying pitch, roll and yaw angles (as defined in figure 5).

When looking at the distributions in figure 6 it is clear that the roll average is 0 (i.e. looking horizontally), and that there is a minor shift (3.16 degrees) towards higher pitch angles (i.e., slightly looking down). The yaw distribution shows a strong bimodal distribution, with one peak of faces depicted with their right cheek and one peak of faces depicted with their left cheek. From the literature, one might assume that the right peak of the yaw distribution, i.e. depicting the left cheek, might be dominated by paintings of females. To test this, we split the pitch, roll and yaw data into a male and female group, based on the Face++ gender predictions. Results are shown in figure 7. Here we can see two gender effects. First, the expected effect on the yaw angles: females generally are depicted showing their left cheek, while males are more likely to be depicted showing their right cheek. Furthermore, we found an interesting novel finding in the pitch distribution: female faces on average displayed a somewhat larger pitch angle than the male faces, meaning that female faces were tended to be depicted to look slightly more down than males. While relatively small (0.82 degrees) the effect is highly significant with a Mann-whitney test ($U=1.48 \times 10^2$, $P<.0001$). To further investigate the relationships between the yaw peaks and the other two orientation characteristics we plotted combined histograms, as shown in figure 8. As there is a clear asymmetry, we fitted a (second order) polynomial through the data. The regression in the histogram for yaw and pitch shows that pitch is on average positive (which is also evident from the histograms in figure 6) but also shows that the pitch

decreases when the face orients outwards in either direction. Secondly, in the combined histogram for yaw roll we see that faces with a lower yaw angle (i.e., faces showing their right cheek, which was shown to be more likely for males) tend to have a positive roll angle, meaning they tilt to their left. As the yaw angle increases, the roll angle decreases by 15% of the yaw. Taken together, it shows that, faces are depicted to look slightly down, and to look further down when the face orients outwards.

Next, we analyzed if these effects of posture changed over time. To this aim we created density plots (figure 9) of the relative frequencies for pitch, roll and yaw angles as a function of year. At the left of figure 9, pitch appears to remain very stable over time. The middle of figure 9

illustrates a larger variability within the roll angles, which can be seen to slowly decrease over time. Last, at the right of figure 9, for yaw over time, we again see two separate clusters for positive (more likely to be female) and negative (more likely to be males) yaw angles, indicating the tendency of faces to be depicted as looking towards the sides. This effect has already been discussed in previous paragraphs, as well as in the literature. Interestingly, this visualization shows a novel finding: the tendency of faces to be depicted as looking outwards is steadily decreasing over time.




Figure 6. Histograms for the faces for production year, and the pitch, roll and yaw angles in degrees.




Figure 7. Normalized histograms for the faces for production year, and the pitch, roll and yaw angles in degrees., where the faces are split into two groups based on gender.




Figure 8. Combined histograms of the pitch angle (left) and roll angle (right) as a function of the yaw angle. A second order polynomial has been plotted in red. On the left histogram for pitch and yaw we see that faces that are looking outwards tend to look up. On the right histogram, we can see that faces tilting to one side tend to show their opposite cheek




Figure 9. The three angles of rotation expressed as relative frequencies a function of year. Pitch shows little variation over year, while roll display less variance as years progress. Yaw shows that the angle with which faces are looking to the side is decreasing over time.

Color

As previously discussed, we constrained our color analysis to global statistics as a function of gender and time. The global statistics are the mean luminance; the variance of the luminance; and our colorfulness metric. We chose to constrain our analysis in this way to deal with the lack of consistent color management and/or calibrations of the paintings.

First, we looked at the effect of gender on color. In figure 10, it is clear that for female skin, the mean luminance is on average higher when compared to male skin (0.513 for female, and 0.473 for male) and the variance of the luminance is on average lower when compared to male skin (0.0135 for female, and 0.0146 for male). This implies that female skin is depicted slightly lighter and with less luminance variation/contrast. Furthermore, the colorfulness metric shows that female skin is, overall, depicted less colorful (63.6 for female and 68.6 for male). Possibly this is a result of female skin being depicted as smoother relative to male skin. This might be a biological gender difference. Perhaps there is an effect of make-up, but this could a priori also result in higher values of colorfulness (e.g. red cheeks with differently hued skin).

Next, we looked at the effect of time over the use of our color metrics. To do this, we chose to further constraint this analysis to three time periods, as the data is not homogeneously distributed throughout history as can be seen in the first histogram in figure 6. There are two clear peaks within the data around 1500 and around 1650. There is also a general increase in paintings between the years 1700 and 1900. The existence of these peaks in our data might be due to a general increase of painting creation, or by museums being more interested in these time periods. These three time periods, 1450-1600, 1600-1700 and 1700-1900 will be used in the analysis of the color data. The first two periods roughly denote Renaissance and Baroque, respectively. The third period covers a number of movements, including neoclassicism, romanticism and impressionism.

We fitted linear models for each of these periods and

plotted them in figure 11. Significant slopes, that is significant change over time, have been colored cyan, while red marked insignificant slopes.

There is a strong decrease in luminance and an increase in luminance variance in the first time period. Next, there is also a very strong decrease in our colorfulness metric in the second time period. However, this appears to be an effect of an even more interesting jump in colorfulness right at the onset of the second time period, which then decreases to an average level during the remainder of this period, only to increase again, but much more gradually, in the 18th and 19th century.

Conclusions

In this study we have created the first, large-scale set of faces from predominantly western art paintings over a large period of time. Using this set, we have shown that painters general do not paint frontal-view faces, instead favoring depictions primarily showing the left or right cheek. Left cheek portraits generally depict females more often, while right cheek portraits generally depict males more often. However, we found that the magnitude of the yaw angles, i.e. the magnitude with which faces are looking outwards is decreasing over time. In other words, these discussed gender differences appear to decrease throughout history. We also found an effect of painted faces to look slightly down, which was found to be stronger for paintings of females. Furthermore, we applied a small proof of concept color analysis on data to which unknown, or possibly, no color calibrations were applied. We found that, overall, variations of the colorfulness and luminance of the painted skin occur throughout history.

Acknowledgments

This work is part of the research program 'Visual communication of material properties' with project number 276-54-001, which is financed by the Netherlands Organization for Scientific Research (NWO).




Figure 10 Three histogram distributions of the mean luminance (left), the variance of the luminance (middle) and the ab surface metric, which is our colorfulness metric. Each histogram has been split up on gender, with male faces represented in grey and female faces in green.




Figure 11. Three scatter plots for the mean luminance, the variance of the luminance and the ab surface measure, with a moving average using a window of 20 years with steps of 2 years (the black line), with 0.5 times the standard deviations up and down, visualized as the grey lines. For three time periods (1450-1600, 1600-1600 and 1700-1900) we have plotted the slope. Blue slopes are significant at a 5% alpha level, while red lines are not significant.

References

- [1] Young, A. W., & Burton, A. M. (2018). Are We Face Experts? *Trends in Cognitive Sciences*, 22(2), 100–110. <https://doi.org/10.1016/j.tics.2017.11.007>
- [2] Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences*, 4(6), 223–233. [https://doi.org/10.1016/S1364-6613\(00\)01482-0](https://doi.org/10.1016/S1364-6613(00)01482-0)
- [3] Grill-Specto, K., Weinier, K. S., Kay, K., & Gomez, J. (2017). The Functional Neuroanatomy of Human Face Perception. *Annu Rev Vis Sci.*, 3, 1167–1196. <https://doi.org/10.1016/j.physbeh.2017.03.040>
- [4] Saleh, B., & Elgammal, A. (2016). Large-scale Classification of Fine-Art Paintings: Learning The Right Metric on The Right Feature. *International Journal for Digital Art History*, 2.
- [5] Bernhard, M. (2016). Gugelmann Galaxy . An Unexpected Journey through a collection of Schweizer Kleinmeister. *International Journal for Digital Art History*, 2.
- [6] Conesa, J., Bruold-Conesa, C., & Miron, M. (1995). Incidence of the Half-left Profile Pose in Single-subject Portraits. *Perceptual and Motor Skills*, 81(3), 920–922.
- [7] Bruno, N., & Bertamini, M. (2013). Self-Portraits: Smartphones Reveal a Side Bias in Non-Artists. *PLoS ONE*, 8(2), 1–5. <https://doi.org/10.1371/journal.pone.0055141>
- [8] Lindell, A. K. (2017). Consistently showing your best side? Intra-individual consistency in #selfie pose orientation. *Frontiers in Psychology*, 8(FEB), 1–7. <https://doi.org/10.3389/fpsyg.2017.00246>
- [9] Bruce, V., Valentine, T., & Badderly, A. D. (1987). The basis of the 3/4 view effect in face recognition. *Applied Cognitive Psychology*, 1(2), 109–120.
- [10] Harries, M. H., Perrett, D. I., & Lavender, A. (2013). Preferential inspection of views of 3-D model heads. *Perception*, 42(11), 1215–1226. <https://doi.org/10.1068/p200669n>
- [11] Liu, C. H., & Chaudhuri, A. (2002). Reassessing the 3/4 view effect in face recognition. *Visual Cognition*, 83(1), 31–48. <https://doi.org/10.1080/13506280244000050>
- [12] McManus, I. C., & Humphrey, N. K. (1973). Turning the left cheek. *Nature*, 243(June 1973), 271–272. <https://doi.org/10.1038/243271a0>
- [13] Nicholls, M. E. R., Clode, D., Wood, S. J., & Wood, A. G. (1999). Laterality of expression in portraiture: Putting your best cheek forward. *Proceedings of the Royal Society B: Biological Sciences*, 266(1428), 1517–1522. <https://doi.org/10.1098/rspb.1999.0809>
- [14] Powell, W. R., & Schirillo, J. A. (2009). Asymmetrical facial expressions in portraits and hemispheric laterality: A literature review. *Laterality*, 14(6), 545–572. <https://doi.org/10.1080/13576500802680336>
- [15] Tyler, C. W. (1998). Painters centre one eye in portraits. *Nature*, 392(6679), 877–878. <https://doi.org/10.1038/31833>
- [16] Carbon, C. C., & Pastukhov, A. (2018). Reliable top-left light convention starts with Early Renaissance: An extensive approach comprising 10k artworks. *Frontiers in Psychology*, 9(APR), 1–7. <https://doi.org/10.3389/fpsyg.2018.00454>
- [17] Lehmann, A. S. (2008). Fleshing out the body: The “colours of the naked” in workshop practice and art theory, 1400-1600. *Nederlands Kunsthistorisch Jaarboek*, 58, 86–109. <https://doi.org/10.1163/22145966-90000743>
- [18] Guest, S., Mcglone, F., Hopkinson, A., Schendel, Z. a, Blot, K., & Essick, G. (2013). Perceptual and Sensory-Functional Consequences of Skin Care Products. *Journal of Cosmetics, Dermatological Sciences and Applications*, 3(01), 66–78. <https://doi.org/10.4236/jcdsa.2013.31A010>
- [19] Koenderink, J., & Pont, S. (2003). The secret of velvety skin. *Machine Vision and Applications*, 14(4), 260–268.

- [20] Anderson, R. R., & Parrish, J. A. (1981). The optics of human skin. *Journal of Investigative Dermatology*, 77(1), 13–19. <https://doi.org/10.1111/1523-1747.ep12479191>
- [21] Matts, P. J., Fink, B., Grammer, K., & Burquest, M. (2007). Color homogeneity and visual perception of age, health, and attractiveness of female facial skin. *Journal of the American Academy of Dermatology*, 57(6), 977–984. <https://doi.org/10.1016/j.jaad.2007.07.040>
- [22] Stephen, I. D., Coetzee, V., & Perrett, D. I. (2011). Carotenoid and melanin pigment coloration affect perceived human health. *Evolution and Human Behavior*, 32(3), 216–227. <https://doi.org/10.1016/j.evolhumbehav.2010.09.003>
- [23] Fink, B., Grammer, K., & Matts, P. J. (2006). Visible skin color distribution plays a role in the perception of age, attractiveness, and health in female faces. *Evolution and Human Behavior*, 27(6), 433–442. <https://doi.org/10.1016/j.evolhumbehav.2006.08.007>
- [24] Angelopoulou, E. (2003). Understanding the color of human skin. *Human Vision and Electronic Imaging VI*, 4299(June 2001), 243–251. <https://doi.org/10.1117/12.429495>
- [25] d'Eon, E., Luebke, D., & Enderton, E. (2007). A system for efficient rendering of human skin. *Proceedings of the 18th Eurographics Conference on Rendering Techniques.*, 147–157. <https://doi.org/10.1145/1278780.1278809>
- [26] Borshukov, G., & Lewis, J. P. (2005). Realistic human face rendering for “The Matrix Reloaded.” *ACM Siggraph 2005 Courses*, 13. <https://doi.org/10.1145/1198555.1198593>
- [27] Igarashi, T., Nishino, K., & Nayar, S. K. (2007). The Appearance of Human Skin: A Survey. In *Foundations and Trends® in Computer Graphics and Vision* (Vol. 3). <https://doi.org/10.1561/06000000013>
- [28] Di Cicco, F., Wijntjes, M. W. A., & Pont, S. C. (2019). Understanding gloss perception through the lens of art: Combining perception, image analysis, and painting recipes of 17th century painted grapes. *Journal of Vision*, 19(3), 1–15. <https://doi.org/10.1167/19.3.7>

Gallery Websites and data link

The individual galleries from which the paintings were downloaded:

1. <https://www.rijksmuseum.nl/en>
2. <https://www.museodelprado.es/en>
3. <https://www.nga.gov/>
4. <https://www.nationalmuseum.se/en/>
5. <https://www.getty.edu/museum/>
6. <https://www.metmuseum.org/>

The data can be downloaded at

<https://doi.org/10.4121/uuid:3beee8ef-1b7e-451f-966f-13230cb2bbe7>

Author Biography

Mitchell van Zuijlen received his B.S. (2016) in cognitive neurobiological psychology and his M.S. (2016) in applied cognitive psychology from Utrecht University, Netherlands, where he focused on perception. Since 2017 he has been a member of the Perceptual Intelligence Lab at the Technical University of Delft, where he is doing his Ph.D.



JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!



- **SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •**
- **INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •**

www.electronicimaging.org

