



Caricature Generator: The Dynamic Exaggeration of Faces by Computer

Susan E. Brennan

Leonardo, Vol. 18, No. 3. (1985), pp. 170-178.

Stable URL:

<http://links.jstor.org/sici?sici=0024-094X%281985%2918%3A3%3C170%3ACGTDEO%3E2.0.CO%3B2-U>

Leonardo is currently published by The MIT Press.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/mitpress.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Caricature Generator: The Dynamic Exaggeration of Faces by Computer

Susan E. Brennan

Abstract—The author has researched and developed a theory of computation for caricature and has implemented this theory as an interactive computer graphics program. The Caricature Generator program is used to create caricatures by amplifying the differences between the face to be caricatured and a comparison face. This continuous, parallel amplification of facial features on the computer screen simulates the visualization process in the imagination of the caricaturist. The result is a recognizable, animated caricature, generated by computer and mediated by an individual who may or may not have facility for drawing, but who, like most human beings, is expert at visualizing and recognizing faces.




Fig. 1. Traditional caricature, drawn by the author, pen and ink, 1981. The act of drawing sometimes hampers the visualization process. (a) A digitized photograph of the subject. (b,c) Two intermediate sketches on the way to a caricature, both of which resemble the subject very little. Each attempt concentrates on only a couple of isolated aspects of the face; these attempts represent dead ends in the visualization process. (d) The finished hand-drawn caricature.

I. INTRODUCTION

However regular we may imagine a face to be, however harmonious its lines and supple its movements, their adjustment is never altogether perfect: there will always be discovered the signs of some impending bias, the vague suggestion of a possible grimace, in short, some favourite distortion towards which nature seems to be particularly inclined. The art of the caricaturist consists in detecting this, at times, imperceptible tendency, and in rendering it visible to all eyes by magnifying it. He makes his models grimace, as they would do themselves if they went to the end of their tether.

Henri Bergson, 1900 [1]

Caricature is a graphical coding of facial features that seeks, paradoxically, to be

more like a face than the face itself. It is a transformation which amplifies perceptually significant information while reducing less relevant details. The resulting distortion satisfies the beholder's mental model of what is unique about a particular face. Caricature, traditionally executed with few lines and loaded with symbols, can be considered a sophisticated form of semantic bandwidth compression.

What goes on in the mind's eye of the caricaturist as she or he exaggerates a face? To what extent does the ability to caricature depend on technical drawing facility, and to what extent does it depend on the powers of observation and the critical filters of memory to capture and magnify the essence of a face? Can these visualization and transformation processes be animated using a computer?

The objectives of my research have been to investigate the traditional activity of caricaturing, to develop a computa-

tional theory for transforming a model of a face into a caricature, and to implement the caricature algorithm on a computer. The Caricature Generator is a software program that provides unique opportunities for the computer graphics image-maker.

II. THE PERCEPTUAL SIGNIFICANCE OF FACES AND CARICATURES

Human faces are such compelling patterns that we see them everywhere—in trees, rock formations or other natural phenomena—whenever such an interpretation is even remotely possible. This fundamental human tendency to interpret an abstract configuration as a face makes possible the compression of the facial image into a very few lines that are, nevertheless, recognizable. This process of abstraction is the beginning of caricature.

Susan E. Brennan (computer researcher), Hewlett-Packard Labs., 1501 Page Mill Road, Palo Alto, CA 94304, U.S.A.

Received 19 September 1984.

A particular human face is a visual pattern which we are able to differentiate easily from thousands of other faces that may be metrically very similar. The psychological literature suggests that there is a complex interaction of perceptual and cognitive stages in face recognition and memory. Pictures of individual faces are difficult to recognize when presented upside down [2, 3] or in photographic negative [4], even though the amount of information in the image is the same as in a face presented right side up. The perception of individuating features can be masked by one strong feature; unusual characteristics such as scars can disguise other features. Young children have been shown to pay attention to more transient characteristics like facial expressions and hats [5] and tend to use these superficial aspects to distinguish among unfamiliar faces. It is likely that a calibration with respect to other faces within an immediate population or shared context takes place when people look at a face; people initially have difficulty distinguishing among members of an unfamiliar race [6]. James Gibson introduced the idea that visual perception distills and encodes key features or 'formless invariants', regardless of the point of view or style of representation of an image. With respect to faces he noted:

In observing a caricature or a political cartoon one often does not notice the lines as such ... but only the information they convey about the distinctive features of the

person caricatured. The caricature may be a poor projection of his face but good information about it. The form of the face is distorted but not the essential features of the face [7].

The associative context in which one sees an image, especially a face, is also important. Eleanor Gibson [8] reported a study in which two groups of subjects were shown a set of scrawls and asked to remember them. One group was told that they would see secret writing and was presented with the scrawls in horizontal orientation; the other group was told that they would see faces and was shown the same scrawls rotated 90 degrees. Subjects were significantly more successful in remembering these patterns as faces than as writing—an intriguing result, since writing is another pattern with great semantic importance to human beings.

In face-to-face communication, a human face can be considered a display of the highest resolution. People are adept at separating the permanent structure of the face from the temporary interplay of expressive musculature due to emotion, speech and aging. We are so sensitive to facial proportions and dynamics that the slightest change in the image of a face may radically alter our perception of its identity or message.

III. THE HEURISTIC METHODS OF THE ARTIST

In developing a theory and an algorithm for caricature, I have used

artists as informants. Leonardo da Vinci [9] and Albrecht Durer [10] were obsessed with extremes of ugliness as well as with ideals of beauty and sketched many variations of the human face. There is no evidence that most of their sketches were meant to represent specific individuals, so, in the strictest sense, they do not fall within the realm of portrait caricature. Leonardo advised artists to observe and remember four principle variations in the profile; he provided pages of these variations, such as noses. This piecemeal approach treated the face as a series of primitives. Durer's variations, on the other hand, illustrated how a rectilinear coordinate system applied to an 'ideal' face could be transformed into the rectilinear coordinates of an idiosyncratic face.

Francis Grose in *Rules for Drawing Caricatures* [11] introduced the idea that caricatures of individual faces actually start with and deviate from some norm. He cautioned that a modest amount of deviation causes laughter, while a great amount of deviation incites horror.

Portrait artists [12] and cartoonists [13, 14] suggest starting with generalized anatomical models for facial proportions. Many caricaturists [15] keep a file of photographs of public figures from which to work and always work from more than one picture. Another method is to study several photos of the subject and then to draw from memory [16].

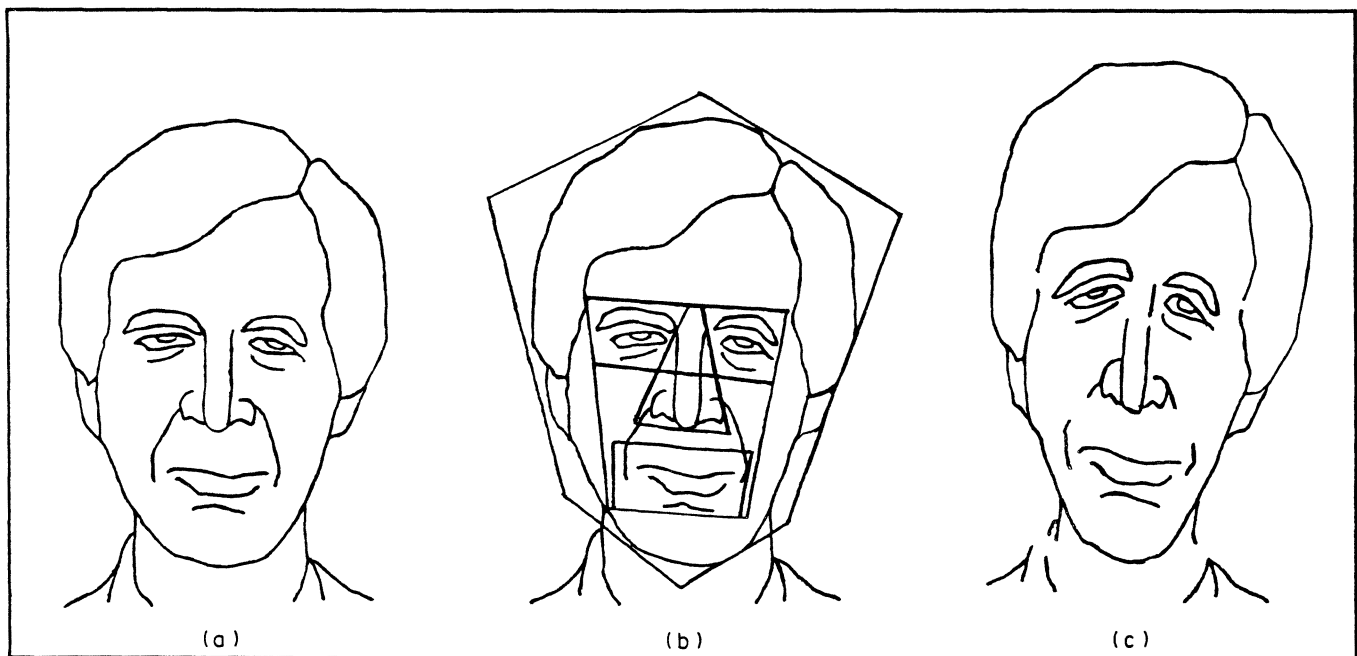


Fig. 2. A computer-assisted caricature made by stretching and squashing areas in a line drawing on a frame buffer; by the author, 1981. (a) The line drawing, traced from the digitized image in Fig. 1a using a digitizing tablet. (b) A face-shaped grid was then superimposed over the line drawing and a collaging program was run which allowed the face to be distorted, step by step. (c) The result of only five steps: lengthen face, lengthen whole head, rotate eyes, widen mouth, and warp whole head by widening at forehead and narrowing at jaw.

In my own experience as a caricaturist, it has been helpful to have a three-dimensional mental model of the subject, either from studying several pictures or from memory of a live model. Caricature is a projection of three-dimensional information over time that contains more information about facial volumes and dynamics than would a simple two-dimensional projection. Ironically, drawing a caricature from a live model can be extremely difficult because of the tendency to record too many details and produce a realistic sketch, as the filters provided by memory are circumvented. Before drawing a caricature I find myself looking away from the subject, closing my eyes and visualizing; I see the whole face all at once in my mind's eye without analyzing any one part of it, and then I watch it amplify itself. However, the tendency while translating memory into drawing is to exaggerate just a few features at a time until the fleeting vision is approximated. The technicality of having to draw one line and then another distracts the caricaturist from the emerging caricature. Figure 1 shows a subject and the sketches leading to my hand-drawn caricature of him. The intermediate stages capture only a few things about the face at a time and contain lines that actually contradict and inhibit the visualization process. This experience corroborates Perkins' study of key facial features [17] which used as stimuli caricatures of Richard Nixon, each of which either left out one of four key parts of the face or replaced it with an incorrect component. His conclusion: contraindicating any one of the key attributes made the caricature unrecognizable, whereas omitting the attribute entirely was not so harmful to the caricature's identity.

When I first began using computer graphics as a medium for caricaturing, I sought to liberate myself from the traditional constraints of making serial, static marks with paper and pencil [18]. Using the same subject shown in Fig. 1, the caricature in Fig. 2 was done in a more holistic fashion by interactively collaging pieces of a line drawing of a face on a frame buffer (Fig. 2a) using a face-shaped grid (Fig. 2b), in five steps: lengthen face, lengthen whole head, rotate eyes, widen mouth and, finally, warp whole face by widening forehead and narrowing jaw (Fig. 2c). I found the drawing in Fig. 2c to be a better representation of the subject than the traced, more 'accurate', line drawing with which I began. But what I really wanted was an algorithm that would enable me to exaggerate a whole face dynamically and in parallel, as I do in my imagination. This fantasy led me to write the Caricature Generator software, which I used to produce the drawings of my subject shown in Fig. 3.

IV. A THEORY OF CARICATURE

It is not really the perception of likeness for which we are originally programmed, but the noticing of unlikeness, the departure from the norm which stands out and sticks in the mind.

E.H. Gombrich [19]

I have made the underlying assumption that a caricature is a thrifty, exaggerated portrait of a specific person. For simplicity's sake I have limited caricature to a black-and-white line drawing of a face. How a caricature is created and how it is perceived may depend on very different cognitive processes; however, both rely on a mental model of a face that can be used to distinguish it from other faces. Unlike some other well-known

computer graphic transformations of abstract faces [20] or typefaces [21], the Caricature Generator does not transform predetermined components using independently varied parameters, but abandons altogether the traditional notion of facial features as components.

Determining key features for caricaturing (and recognizing) a human face is an enormously ambitious undertaking, particularly since there is no evidence that this feature set is the same from context to context, from beholder to beholder, or even from face to face. Deciding how to describe an individual feature can be a difficult and somewhat arbitrary process. There are no absolutes; the dimensions of one area influence the perceived shape of an adjacent area. The 'features' that the English language names are totally inadequate to synthesize a complete facial description. These words tend to correspond instead to the sensory organs—the face's inputs and outputs. It is not only the eye, but its orientation with respect to the rest of the face, that is memorable. Understanding facial features is still an unsolved problem for machine vision.

A few studies have tried to determine experimentally a set of critical features that caricature addresses. Perkins [17] surveyed existing popular caricatures of Richard Nixon and discovered that a small set of features was fairly consistently chosen for exaggeration by newspaper artists. This feature set was not general to all faces, but specific to the face of Richard Nixon. Goldman and Hagen [22] selected 11 'feature ratios' that seemed distinctive to the face of Nixon and included Perkins' four features. They found that political cartoonists were fairly consistent in the choice of which features in Nixon's face to distort, but that there was an enormous variation in the extent to which these features were distorted. Their conclusion was that the degree of distortion was a function of style and political bias.

I have based this work on the hypothesis that a caricaturist selects and amplifies features that make a particular face unique. Clearly, it is on the right track to say of a face, "His nose is unusually long—make it even longer". The problem is in deciding which things to measure and what to compare them with. How can one determine what is 'unique' about a face? Implicit in some of the heuristic methods described earlier is a norm for comparison, although it is unlikely that caricaturists apply such a standard consistently or even consciously. By making measurements of many faces, one could compile a generic face that



Fig. 3. Drawings by the Caricature Generator of the same subject as in Figs 1 and 2, 1982. (a) Four stages, from realistic line drawing to extreme caricature, superimposed. The continuous exaggeration in the progressive images is evident. (b) Caricature of the subject with respect to an average-white-male composite face.

approximated an average of all the faces an individual has ever seen; alternatively, based on the variation within a population of faces, one could determine the most average existing face [23]. One could try to draw the most nondescript face possible to visualize. One might choose the published ideal of an anatomist or aesthetician. Or one might use one's own face as a basis for comparison.

The theory of computation underlying the Caricature Generator is to exaggerate the metric differences between a graphic representation of a subject face and some other similarly structured face, ideal or norm. This norm is meant to correspond to the hypothetical model of a generic or average human face which resides in the mind's eye of the artist and provides a basis for judging what is unique about a face. The critical process of selecting what to include and what to leave out (which

comes so naturally to the human caricaturist) is finessed by having the computer exaggerate *all* spatial relationships and by delegating responsibility for choosing the basis for comparison to the human user of the system. Thus the system makes no qualitative decisions about individual distinctive features.

Implicit in the theory behind the Caricature Generator, then, is the convenient notion that a relationship between lines on the subject face becomes a 'feature' only when it differs significantly from the corresponding relationship on a comparison face—in other words, when it becomes useful in distinguishing one face from another.

V. COMPUTER-GENERATED VS TRADITIONAL CARICATURE

It is not the intent of this work to replace the human artist, but merely to

provide a way for skilled and unskilled caricaturists to focus on aspects of the process other than the actual drawing. There are, as one would expect, major differences between hand-drawn and machine-generated caricatures.

Whereas a human cartoonist leaves out those elements determined insignificant, the Caricature Generator merely exaggerates less where the differences are less. Perkins [17] aptly delineates several theories of caricature, among which is a 'selection' theory which says that when people look at caricatures, they not only notice key attributes but also ignore the negation or absence of certain details. Therefore, as long as the most important spatial relationships within the subject face are amplified, the distortion of details should not cause the caricature to be unrecognizable, but should merely make it a bit more 'noisy'—an effect that could even be interpreted as an element of style.

Another departure from traditional caricature is that lines are not constrained to stay connected to or distinct from one another, which means that an eye is free to float above an eyebrow. An early, more 'intelligent' prototype of the Caricature Generator sought to constrain some of these possibilities. But I discovered that users enjoyed having greater control over the developing image—for example, exploring how far an image could be exaggerated before becoming unrecognizable or drifting into 'facelessness'.

There are, of course, other important differences between caricatures generated by hand and by machine. The Caricature Generator currently makes portrait caricatures without incorporating any of the political or other contextual symbols and transformations that are an entertaining element in traditional caricatures. Encoded in traditional caricature is not only the identity of the subject but also that of the artist, as revealed by the drawing's style. When viewers encounter a caricature, they interpret the style and other variables such as political context and expressiveness. Needless to say, the stylistic elements of caricature are still beyond a computer's grasp. Such things as line quality, number of points used to represent the face, degree of distortion and choice of ideal can be systematically varied to approximate some of the elements of style. The area of representational style is a promising topic for further research in artificial intelligence.

VI. IMPLEMENTATION

The Caricature Generator has been implemented as a computer graphics




Fig. 4. The user interface of the Caricature Generator, printed directly from a Symbolics Lisp Machine screen by the author, 1984. The caricature window is at the top of the screen and the exaggeration control menu is below. The user moves the cursor (here, in the shape of a black arrow) with a mouse input device and selects a point along the black horizontal scale beneath the two icons. Choosing a point directly beneath an icon causes the upper window to display what looks like a traced drawing of that face. Points farther out on the scale yield caricatures of the closer icon with respect to the farther one. Points on the scale between the two faces yield average, or 'offspring' faces. The user may also move the cursor into the caricature window, grab a point on a line and stretch it directly, as a 'rubber-band' face.

process that compares two faces and then synthesizes a caricature. Each face is schematically represented as a line drawing, which is an object consisting of an identity and a set of lines. Each line is an object which consists of a name and a set of points through which passes a smooth curve. These points for any given line are consistent in number and order for each face in the database. Faces lacking wrinkles or moustaches contain virtual lines which remain invisible until needed as a basis for comparison. Initially the two faces chosen by the user are normalized: that is, they are scaled and translated relative to one another by spatially aligning the pupils of their eyes. During the comparison, points from the subject face are mapped onto corresponding points on the norm or comparison face. The distance between each pair of corresponding points on the two faces is represented as vector. Each vector is then multiplied by an amount of exaggeration selected by the user. A caricatured line drawing of the subject relative to the norm is rapidly generated using the set of points which results from adding the exaggerated vectors to the points of the subject face. This process can be more easily understood as the converse of the animation technique of in-betweening: rather than averaging points together, the distance between them is increased.

The display consists of two windows: one in which the caricature appears, and one containing the exaggeration control menu with which the user controls the

process (Fig. 4). The exaggeration control is a horizontal linear scale upon which are displayed black and white (one-bit) icons of the two faces being compared. By choosing a point directly below an icon, the user causes the unadulterated line drawing of that face to appear in the caricature window above. By choosing a point on the end of the scale beyond the subject face, the exaggeration is increased so that the result is a caricature of the subject with respect to the comparison face. By choosing a point on the opposite end of the scale beyond the comparison face (i.e. by turning the exaggeration down into the negative or left-hand region of the scale) a caricature of that face with respect to the subject face is generated. By choosing a point between the two faces, an averaged or idealized face results. The menu allows the user to see the subject and comparison faces simultaneously and to visualize the complex changes she or he is initiating on the caricatured face. The user's mental model of the process is therefore a spatial and dynamic one; the user is making one face look less like another than it does already. The reversibility of the continuous transformation drives home this relationship.

At present the basic line drawings that make up the database have been input in advance by an experienced user of the system. Fewer than 200 points are rotoscoped over a digitized image of a frontal face, using a mouse input device. As each point on a particular line is

drawn, an interpolating curve [24] is generated which connects the points. The choice of which lines on a face to include in the representation was made with the goal of presenting a recognizable face using the fewest number of lines possible. These lines are intended to include those most often used in caricatures drawn by human artists and correspond roughly to such perceptual features as occluding edges, material changes (as from hair to skin) and dark/light transitions on the face. Figure 5 illustrates the line representation that results from the input program.

Once a face is stored as a line drawing it is available to be interactively distorted and animated by the casual caricaturist. No technical drawing skill is required. Figure 6 illustrates the progressive exaggeration of the face of John F. Kennedy.

Originally the basis for comparison used in the Caricature Generator were such faces as the average of all the other faces in the database, a norm created from multiple exposures of a population of white males from Aspen, Colorado [25], some anatomical ideals after artists Leonardo da Vinci [9], Oskar Schlemmer [26] and Albrecht Durer [10], and my own face. In the process of creating a large database that includes both famous faces and personally familiar faces, I have discovered that frequently a successful caricature results from using as a norm any face that just seems very different from the subject face. One may caricature face A with respect to face B, and face B with respect to face A, as illustrated by the sequences in Fig. 7. The selection of a basis for comparison is entirely subjective and indeed appears to contribute significantly to the amusing discoveries reported by users of the system. It also throws into question the idea that there need be only one strong norm for all human faces.

It is significant that the amount of data stored for each face in the database is only on the order of 400 bytes per face. This thrifty, object-oriented description of a face is easy not only to manipulate and store but to animate as well. A specialized animation package for the Caricature Generator automatically generates cycles of dynamic facial gestures, such as expressions and lip movements, for any particular caricature. This animation process is heuristically consistent with the rest of the caricaturing process. Each frame in an animation cycle is created by comparing each line to the corresponding line in a stored template cycle (for, say, a generic smirk). Cycles are generated and stored in advance and can be displayed

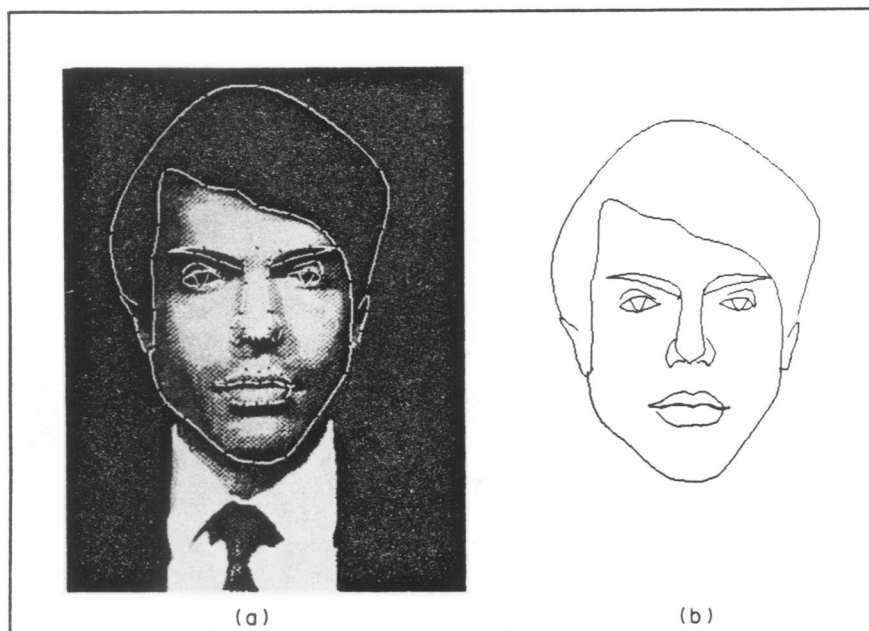


Fig. 5. Input to the Caricature Generator by the author with the Caricature Generator, 1982. (a) A digitized photograph on top of which a hundred or so points were entered using a typical input device (tablet or mouse). The points were then connected with a smooth curve by the computer. (b) The resulting line representation. Once a face is input in this way, it is available to be manipulated in a variety of ways by any user.

rapidly in any order on a frame buffer. In this way any caricature can be made to talk and grimace automatically and interactively. The animated behavior can be controlled by other programs or correlated with external events such as sounds generated by a text-to-speech synthesizer or commands typed on a keyboard. An object-oriented, bandwidth-limited description of a face is a plus for computer graphics and animation.

The original implementation of the Caricature Generator was programmed in the PL/1 language on a mainframe computer and used a digitizing tablet and a frame buffer with a touch-sensitive surface. The current implementation is in the Lisp language on a Symbolics machine [27] which uses a mouse for input. The window system, pop-up menus, high-resolution screen and other features in this programming environment lend themselves conveniently to the user interface.

VII. ADDITIONAL APPLICATIONS AND DIRECTIONS

Caricature can be considered a form of semantic bandwidth compression. This term refers to the concept of representing or transmitting only that information or part of an image that is most meaningful in a particular context. My research in caricature began as part of a teleconferencing project [28]. The goal was to represent and transmit over limited bandwidth (such as a telephone line) some of the visual nuances present in face-to-face communication. It was discovered that animated caricatures of faces were more acceptable in a teleconferencing situation than were some of the more realistic synthesized images of talking heads, because the caricatures made the

degree of abstraction in the image more explicit.

Users report that the Caricature Generator is an effective tool for exploring the very subjective, qualitative character of an individual face and for demonstrating how far one can exaggerate and still recognize a face. It works because the user's mental model of the process is consistent with the transformation of the face on the screen. The Caricature Generator is amusing because it can exaggerate a face with respect to another extremely different face which the user selects simply by browsing through the database of faces. It is intriguing because it sheds some light on an elusive imaginative process. It is satisfying also to the 'Sunday painter' [29]; a person without traditional artistic skill can manipulate a visual image in a semantically complex way.

I have also explored techniques for fully automating the process of adding new faces to the database. Particularly promising are those predictive schemes found in the machine vision literature concerning face recognition algorithms which find points by knowing where to look and by relying on special knowledge about the symmetry and anatomy of the human face [30, 31]. If the input stage can be fully automated, then the Caricature Generator will be able to exaggerate and animate any face that pauses in front of a digitizing camera. This capability will make possible entirely new applications, such as the ability to digitize the user's face and immediately place it as an animated caricature into a computer game, workspace, mail message or interactive story.

Many aspects of caricature remain to be explored. The intriguing variable of representational style has already been

mentioned. The Caricature Generator could be used to analyze drawings of the same face by different artists in an attempt to isolate definitive elements in their individual styles. Since a good caricature contains more information than a simple two-dimensional projection, it would also be promising to explore the use of three-dimensional input. A convention in traditional caricaturing is to draw a three-quarter-view nose on a frontal face; the Caricature Generator could make use of such conventions for combining profile with frontal exaggerations.

The general concept of caricature as an amplification of that which distinguishes one thing from another within a population of similar things could be applied to phenomena other than faces, such as other visual patterns and the articulation of gestures and motion in animation. The challenge is to find an appropriate representation and basis for comparison that will yield perceptually interesting results when caricatured.

It is conceivable that a user would want to create an interactive persona on a computer screen by editing a face, a voice, gestures and, eventually, behavior and style. This process could be done by deliberately using the Caricature Generator as a tool or, if there is sufficient intelligence in the system, through a qualitative dialogue between the user and some stereotyped presence that acts as guide and point of departure on the way to an interactive character. The Caricature Generator could provide the graphical representation for this agent in the machine.

VIII. CONCLUSION

This approach to drawing provides several unique opportunities for com-




Fig. 6. Progressive exaggeration of the face of John F. Kennedy by the author with the Caricature Generator, 1984. The author finds the continuous display of a sequence such as this to be a good model of what goes on in her mind's eye while caricaturing. (a) Undistorted line drawing; no exaggeration. (b) 50% exaggeration with respect to an 'average' face. (c) 100% exaggeration, the author's choice as the 'best' caricature in this sequence. (d) 140% exaggeration. (e) 160% exaggeration.

puter graphics image-makers, who today tend to be programmers, artists or artist-programmers and tomorrow may very well include most computer users.

Many computer graphics tools require a great deal of dexterity and image-making experience on the part of the user. Many graphics interfaces demand that an image be built up from primitives. Most rely on the user's ability to switch from one mode to another or to make correspondences between the visual domain and the tactile or textual domains. The activity of creating images on a computer often feels like a particularly cumbersome combination of doing math and drawing with one's feet.

The Caricature Generator differs from most computer graphics drawing and image-processing systems in that it enables the manipulation of a complex set of spatial relationships in a very intuitive way. It facilitates a visual dialogue between a computer-generated image and one's mental model of a face. The Caricature Generator allows the user to concentrate on the visualization process rather than on the act of drawing with a computer. The user need not employ traditional artistic skills or analytic mental processes necessary to translate a description of a human face into words or brushstrokes. However, the Caricature Generator does not preclude use of these

abilities. Figure 8 is the handiwork of one person who used the Caricature Generator up to a certain point and then decided to take more responsibility for the image. By moving the cursor from the exaggeration control menu up into the caricaturing window, one can directly grab and stretch individual lines as one would a rubber-band face.

Finally, one of the most significant aspects of the Caricature Generator is the way in which the user interface presents the process. Using the system feels remarkably qualitative because of the spatially orchestrated selection process and the power to express computationally and precisely such formerly vague inten-

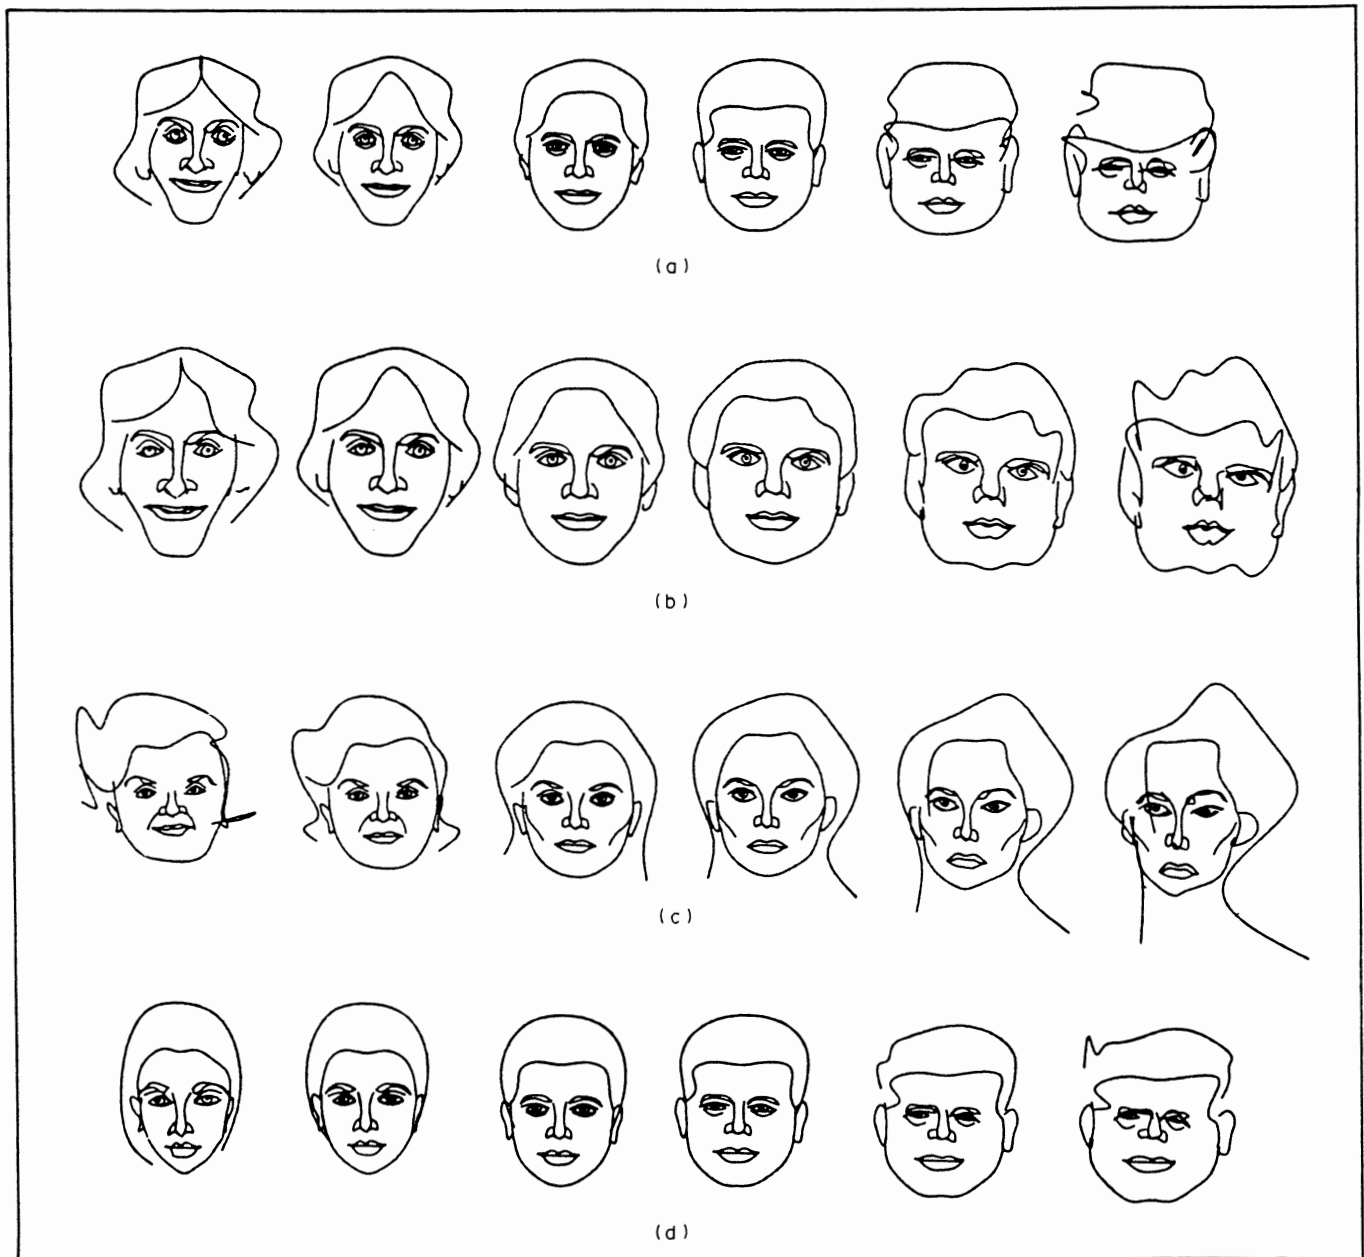


Fig. 7. Continuously exaggerated sequences of pairs of faces. In each sequence, the second and fourth faces are the undistorted faces, the third face is the average (idealized or offspring) face, the first and fifth faces have been exaggerated 50% in their respective directions, and the sixth face is an extreme caricature of the fourth face with respect to the second; by the author with the Caricature Generator, 1984. (a) Craig Reynolds and John F. Kennedy. (b) Craig Reynolds and Howard Cannon. (c) Dianne Feinstein and Fay Dunaway. (d) Elizabeth Taylor and John F. Kennedy.




Fig. 8. Interactively warped (as opposed to caricatured) 'rubber-band' faces, same subject as in Fig. 5; by an anonymous user and the Caricature Generator, 1982.

tions as "make him look less like Nixon", "exaggerate the face even more", and "now try making her look less average". In its current configuration the Caricature Generator is an example of a successful creative partnership between human and machine, where each is allowed to do what it does best.

Acknowledgements—My initial research on this topic was supported by the Cybernetics Technology Division of the Defense Advanced Research Projects Agency and was done at MIT's Architecture Machine Group with thesis supervision and encouragement from Nicholas Negroponte. Subsequent research was done at Atari, Inc. Thanks also to Alan Kay and James Cutting for their intellectual support and to Gary Phipps for digitizing software and hardware.

REFERENCES AND NOTES

1. Henri Bergson, "Laughter" (1900), from *Comedy*, Wylie Sypher, ed. (Baltimore and London: Johns Hopkins University Press, 1956).
2. J. Hochberg and R. E. Galper, "Recognition of Faces: An Exploratory Study", *Psychonomic Science* 9, 619-620 (1967).
3. R. Yin, "Looking at Upside-Down Faces", *Journal of Experimental Psychology* 81, 141-145 (1969).
4. R. E. Galper, "Recognition of Faces in Photographic Negative", *Psychonomic Science* 19, 207-208 (1970).
5. S. Carey and R. Diamond, "From Piecemeal to Configurational Representation of Faces", *Science* 195, 312-314 (1977).
6. R. E. Galper, "Functional Race Membership and Recognition of Faces", *Perceptual Motor Skills* 37, 455-462 (1973).
7. J. J. Gibson, "On the Concept of Formless Invariants in Visual Perception", *Leonardo* 6, 43 (1973).
8. Eleanor J. Gibson, *Principles of Perceptual Learning and Development* (New York: Appleton-Century-Crofts, 1967).
9. Emery Kelen, ed., *Leonardo da Vinci's Advice to Artists* (Nashville, TN: Thomas Nelson, 1974).
10. Albrecht Durer, *Proportionslehre* (Unterschneidheim: Verlag Walter Uhl, 1969; reprinted from 1528).
11. Francis Grose, "Rules for Drawing Caricatures", bound with William Hogarth's *The Analysis of Beauty* (printed in the Strand: for Samuel Bagster, 1791).
12. William Rimmer, *Art Anatomy* (New York: Dover, 1961; reprinted from Little, Brown, & Company, 1877).
13. Jack Hamm, *Cartooning the Head and Figure* (New York: Grosset & Dunlap, 1967).
14. John Adkins Richardson, *The Complete Book of Cartooning* (Englewood Cliffs, NJ: Prentice-Hall, 1977).
15. Dick Spencer, *Editorial Cartooning* (Ames, Iowa: The Iowa State College Press, 1949).
16. Roy Paul Nelson, *Cartooning* (Chicago: Contemporary Books, 1975).
17. D. N. Perkins, "Caricature and Recognition", *Studies in the Anthropology of Visual Communication* 2(1) (Spring 1975).
18. Susan E. Brennan, "Caricature Generator" (Cambridge, MA: Massachusetts Institute of Technology, unpublished thesis, September 1982).
19. E. H. Gombrich, *Art and Illusion* (Princeton, NJ: Princeton University Press, 1960).
20. H. Chernoff, "The Use of Faces to Represent Points in K-Dimensional Space Graphically", *Journal of the American Statistical Association* 68, 361 (1973).
21. Donald E. Knuth, "The Concept of a Meta-Font", *Visible Language* 16(1), 3-27 (Winter 1982).
22. M. Goldman and M. A. Hagan, "The Forms of Caricature: Physiognomy and Political Bias", *Studies in the Anthropology of Visual Communication* 6, 30-36 (1978).
23. Leon D. Harmon, "The Recognition of Faces", *Scientific American*, pp. 70-84 (November, 1973).
24. David F. Rogers and J. Alan Adams, *Mathematical Elements for Computer Graphics* (New York: McGraw Hill, 1976).
25. P. Weil, "About Face", (Cambridge, MA: Massachusetts Institute of Technology, unpublished thesis, September 1982).
26. Oskar Schlemmer, *Man* (Cambridge, MA: M.I.T. Press, 1971).
27. Daniel Weinreb and David Moon, *Lisp Machine Manual* (Cambridge, MA: Symbolics, July 1981).
28. This project took place at Massachusetts Institute of Technology's Architecture Machine Group. See N. Negroponte, A. Lippman, and R. Bolt, "Transmission of Presence" (Cambridge, MA: Architecture Machine Group proposal to the Cybernetics Technology Division of Defense Advanced Research Projects Agency, M.I.T., 1980).
29. N. Negroponte, "The Return of the Sunday Painter", from *The Computer Age: A Twenty-Year View*, M. L. Dertouzos and J. Moses, eds. (Cambridge, MA and London: MIT Press, 1980).
30. Michael D. Kelly, "Visual Identification of People by Computer", Stanford Artificial Intelligence Project Memo AI-130 (Stanford, CA: Computer Science Department, Stanford University, 1970).
31. Takeo Kanade, "Picture Processing System by Computer Complex and Recognition of Human Faces" (Kyoto, Japan: Department of Information Science, Kyoto University, thesis, November 1973).

GLOSSARY

artificial intelligence—the study of ideas which enable computers to appear to behave intelligently.

byte—a word of digital information typically consisting of eight bits, where a bit is either "1" or "0".

caricature—for these purposes, a line-drawing portrait of an individual that exaggerates spatial information about the face with respect to some norm.

database—a structured collection of pieces of information; in this case, a collection of digitized faces.

frame buffer—computer memory that contains an array of values corresponding to the colors of points displayed on a screen.

icon—a graphic representation with some semantic content; a sign.

interpolating spline curve—a smooth continuous curve that passes through a set of knot points.

menu—a presentation of available selections.

mouse—a rolling spatial input device for a computer, frequently used to make a selection on a menu or to move a cursor on a screen. The

mouse is moved in a virtual space that does not bear a constant relation to the geometry of the surface it is on; i.e. if it is picked up and placed elsewhere on a surface, the cursor will not move.

object-oriented—a paradigm for computer programming which builds things out of distinct packages consisting of both information and directions on how to manipulate it.

one-bit—having only one of two possible values (1 or 0); black and white.

primitives—small indivisible pieces; components.

rotoscoping—an animation technique for entering positional information from 'real' imagery (photographs or moving pictures); tracing.

semantic bandwidth compression—representing or transmitting only that part of a message or piece of information that is the most significant in a particular context.

tablet—an input device for a computer that includes a surface and a hand-held, pen-shaped implement. The position of the pen on the tablet surface corresponds to the position of the cursor on the screen.

touch-sensitive display—a computer screen that also acts as a surface for spatial input. A user may move a cursor or make a selection by pointing at or directly touching an item on the screen.

user interface—that part of a system which facilitates the use of a computer by a human being.

window—a rectangular area displayed on a computer screen which provides a discrete environment for manipulating computational objects.