UNIVERSITY OF CALIFORNIA

Santa Barbara

Data-Masks
Biometric Surveillance Masks Evolving in the Gaze of the Technological Other

A Thesis submitted in partial satisfaction of the requirements for the degree

MASTER OF SCIENCE

in

MEDIA ARTS AND TECHNOLOGY

by

Sterling Crispin

Committee in charge:

Marcos Novak, Chair

Lisa Jevbratt

Dr. JoAnn Kuchera-Morin

Dec 2014

The Master's thesis of Sterling Crispin is approved.

Dr. JoAnn Kuchera-Morin, Professor

Lisa Jevbratt, Professor

Marcos Novak, Professor, Committee Chair

Dec 2014

Data-Masks: Biometric Surveillance Masks
Evolving in the Gaze of the Technological Other

by

Sterling Crispin

ABSTRACT

Figure 1:

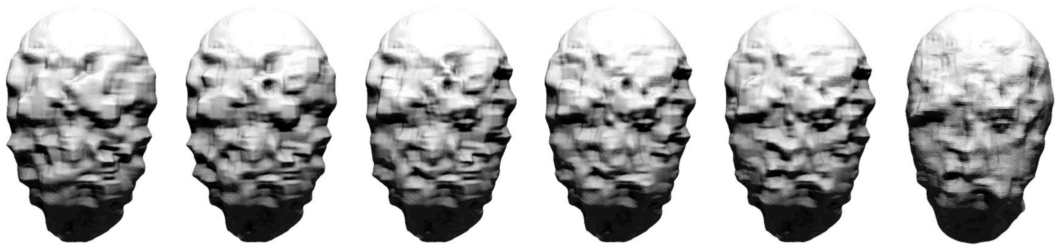

The following work is a theoretical and technical investigation into the form and function of biometric surveillance technology, which is the mathematical analysis of biological data.

Theoretically, I am concerned with the aggressive overdevelopment of surveillance technology, how this is changing human identity, and how humanity interacts with technology. By technology I mean individual instances of technological devices and networked systems like cameras and software, but also what I identify as the 'Technological Other', a global living super-organism of all machines and software. Technically, my specific focus has been in reverse engineering facial recognition, facial detection, and image correlation techniques in order to reveal how they represent human identity.

This research has resulted in the production of a series of 3D printed face masks which were algorithmically evolved to satisfy facial recognition algorithms (Figure 0.1). It is important to understand the goal of creating these masks isn't to defeat facial recognition or provide something undetectable, simply covering your face with your hand

will do that. Rather, my goal is to show the machine what it is looking for, to hold a mirror up to the all-seeing eye of the digital-panopticon we live in and let it stare back into its own mind.

These masks are shadows of human beings as seen by the minds-eye of the machine-organism. They are intended for use in acts of protest, poetry, civil disobedience, and shamanistic ritual by the citizens of our global village as it becomes further blanketed by techno-sphere.

Contents

0.1 General Overview

Generally speaking, Data-masks were created by randomly making images, measuring those images against a facial recognition system, and then mixing the best attempts together to form new attempts. When the images are randomly changed their face-likeness is measured, only good mutations are kept and so this guides them toward a face. This is happening across a population of five or more masks, and the best attempts are recombined and mutated to create new ones in a simple genetic algorithm.

One way I am generating these images is by representing 3D geometry as a volume, randomly filling some of this volume, and then calculating the surface that would encapsulate the generated volume. This approach to volumetric geometry is sometimes used for smoke, water, and fluid special effects in the film industry. The result of this approach is often amorphous fluid-like faces.

Another way I am making masks is by creating 2D images by randomly placing different valued rectangles in them. Once I generate a 2D face I extrude it out of the surface of an averaged 3D head. This allows for more rapid development of the masks and is closely related to how Facebook [1] represents and classifies human identities, which is explored in section 0.2.3.

It is important to understand that facial recognition is not done by simply measuring the distance between your nose and your eye. Modern facial recognition techniques work by abstracting many images of one person into complex mathematical objects. But before you can recognize a particular face you need to detect if there is a face within an image, so you need to build a very general and robust model of what a person is. In this context Data-masks can be understood as visualizations of how machine learning algorithms generalize faces into abstract feature sets.

0.2 Motivations

The creation of these Data-masks is an act of political protest by means of bringing transparency to the surveillance and biometric techniques used today. Data-Masks give form to an otherwise invisible network of control and identification systems and make their effect on our identities tangible and visible.

> "You can't hit what you can't see, you can't grab what you can't touch. You can't critically engage with technoculture and its infrastructure if you're unable to unravel its threads, run your fingers through the seams, visualize its jurisdiction and weigh its influence on everyday life." - Alice Politics summarizing #Stacktivism, a critical social movement which engages digital infrastructure [2].

We live under the shadow of a totalitarian police state and this is how it sees human beings, as abstract things, patterns and numbers, not as individual people whose lives matter. For example, if you have a drivers license in the United States your identification photo is entered into a database with criminal records and searched daily as a potential suspect by the FBI's Next Generation Identification System [3] [4]. This is not how our justice system was designed to work.

Data-masks also existentially question humanity's intimate yet perilous relationship with technology. These masks represent a search for a spirit in the machine by creating false positives and ambiguous identities which delve into the uncanny valley [5], the space between the human and the post-human, and the subconscious of our technological systems. The volume of communication between machines in the industrial internet already far surpasses human to human communication and it is increasing exponentially [6]. We need ways of reaching into this space and retrieving artifacts which give back to the human, and address the human as human. There is a palpable and visceral spark of experience when two people look into each others eyes, and remnants of this essence still exist within these masks.

0.2.1 Biometrics and Surveillance

Biometrics is the measurement and analysis of biological data which can involve biological patterns like fingerprints, facial geometry and patterns, eye retinas and irises, DNA, voice, and body smells, as well as behavioral traits like a person's manner of walking or travel patterns between physical locations [7]. This biometric data is collected and used to represent, identify, and verify an individual's identity. Biometrics can be traced to Charles Darwin's study of the evolution of humans [8], Francis Galton's statistical analysis of human differences [9], and the social philosophy of eugenics which was inspired by their efforts and peaked at the height of WWII with the Nazi movement.

The FBI has developed a system called Next Generation Identification [3] which went fully operational in September of 2014 and has been designed to contain all of these types of biometric data. This data includes the millions of fingerprints already on file, and as many as 52 million faces by 2015 [4]. It is important to note that this database blurs the boundary between data from convicted criminals and innocent citizens which raises important privacy and civil rights concerns. When searching for a person in this system the resulting top 50 matches may only contain the correct person 85 percent of the time [4]. This means that many innocent people will be presented as suspects for crimes in which they had no involvement.

However, it should be obvious that biometric data is not the only personal data being collected and analyzed. The United States and various multinational corporate entities have constructed an Orwellian police state with programs like PRISM [10], the successor to the Total Information Awareness program [11]. As Edward Snowden has revealed, these organizations collect and monitor nearly all forms of communication and collect every piece of digital information online for government based surveillance purposes [12]. Some of the uses of this collected data include the Chicago Police Department's Predictive Policing system which algorithmically compiles a "heat list" of roughly 400 people who are likely to be involved in a crime in the near future [13]. Chicago police have personally visited at least 60 of these individuals prior to them ever committing a crime.

Many private corporations are also creating biometric databases with commercial applications. An example of which is FacialNetwork.com [14], which is a cloud based biometric database that software developers are using to create mobile based recognition software for the public. Apps like theses will soon enable anyone with a mobile device to photograph someone and search the web for their identity and personal information.

The implications these programs have on our privacy, identity, and their impact on social interactions is immense. They suggest a world in which we are always being seen, watched, and analyzed, and new algorithms will emerge ever rapidly which further leverage this data against us.

Data-Masks have been developed in order to make these threats to our identities visible through illustrating the way these networked systems capture, classify, and represent human identity.

0.2.2 Biometrics and the Quantified-Self

If private citizen's personal information, social graphs, and communications are being analyzed then the results should be made available to said persons to empower rather than enslave them. This attitude has become popular in personal fitness, but not in communications, biometric identity, or social networks.

The Quantified Self industry [15] claims that when biometrics and data analysis tools are used to inform things like exercise, diet, health care, and daily routines they can enhance your life. This is a booming industry with hundreds of devices on the market all promising to track your every step and heartbeat, with the holy grail being an artificial intelligence system that would analyze and help users make decisions from this data [16]. Imagine having an AI system that would recommend what foods to eat, and when, based on your peak metabolic rate and nutrient levels.

However, if all this biometric data about your daily activity is sent unencrypted to servers in the cloud, then they become self-surveillance tools. Such is the case with many commonly used products and platforms on the Internet. If you are not paying for a service like Twitter, Facebook, or e-mail, then you are the product, and your data is being sold to the highest bidder. These services collect users communication and behavior online then analyze this data in an attempt to market consumer goods to the users, or predict their behavior. This kind of tracking, analysis, and profiling will enter the physical world as more self-tracking gadgets become commonplace. Devices like the Apple Watch [17] are data collection campaigns disguised as consumer goods. They are literally peering into our bodies and uploading what they see into private corporate servers for analysis and financial gain.

0.2.3 Biometrics, Personal Identity, and Facebook

What are we as humans and as individual people? Are we simply a series of statistical probabilities and biometric patterns? Systems of representing human identities have contained with them many ontological assumptions about what it is to be an individual, and what personal identity is. The previously described surveillance and biometric identification systems define the human as a "what" ie: that which can be measured, not as a "who" ie: our inner self. When human identity is extruded through the instruments of mathematics and computation there is a great reduction of the human.

Figure 2:

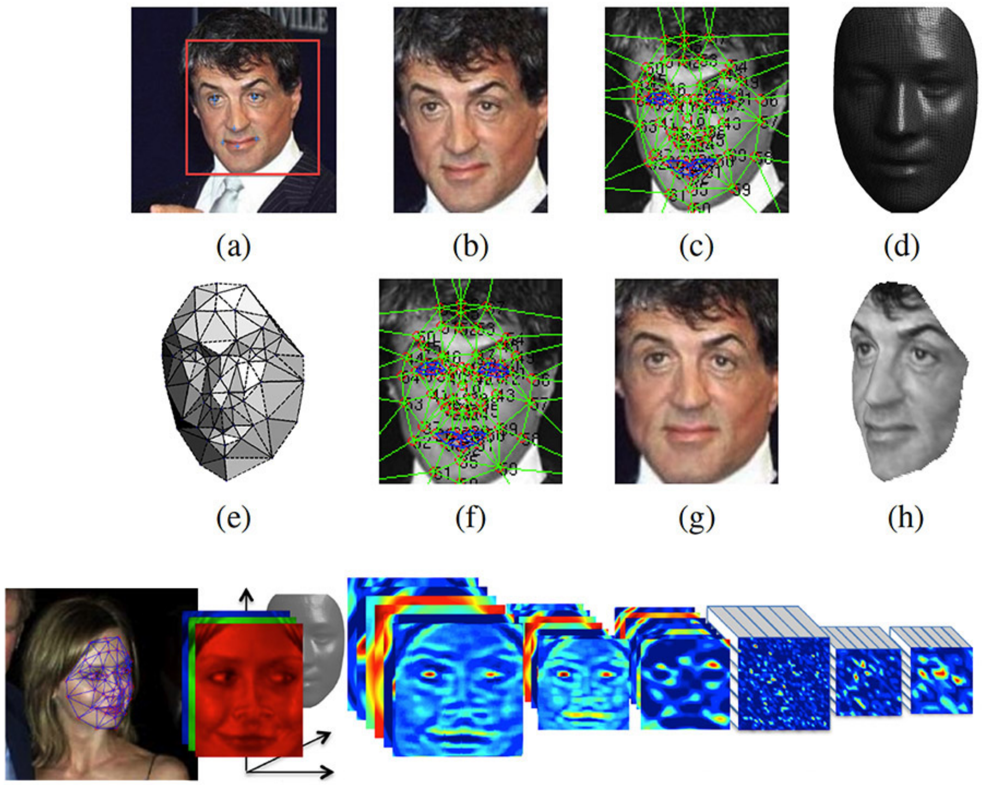

As an example we can look to Facebook's DeepFace [18] system of representing human identities as 3D models and statistical features (Figure 2). Their pipeline involves detecting a face in an image (a), isolating that face (b), measuring 67 "fiducial points" which describe where your facial features are (c), transforming a generic 3D face shape (d) to the image plane of the cropped face, directing a simplified model of the face (e) into the 67 points (f), and finally applying the found face as a texture to the 3D shape so it can be rotated to face front (g) or any other view (h). Seen below these 8 images is an overview of Facebook's DeepFace system. The blue colored feature maps have been created by machine learning algorithms and deep neural nets to represent individual faces in as sparse a way as possible for fast computation and recognition. To put this into further context, neural nets are interconnected systems modeled after how the human brain and nervous systems function, and the Deep

Face system is a nine-layer deep neural network containing more than 120 million parameters [18] describing a face.

This dramatic abstraction and processing of human identity is out of necessity and efficiency for computational reasons, but the results are disturbing. This is in part due to the ghastly disfiguration of the human face and the dismemberment of its features. Additionally, this is disturbing because none of the knowledge gained from this analysis is shared with the individuals analyzed. If the state of the art in computer science can produce a unique feature that describes an individual as such, what good does that do the individual if this knowledge is only leveraged against them? Why are we building an Other for humanity, only to police ourselves with it?

0.2.4 The Rise of the Technological Other

We are witnessing the rise of a globally networked technological organism, the Technological Other, which is a product of emergent coevolution with the human species. It has a distributed body of computers and machines, and a distributed mind of software and algorithms. The Technological Other is the body of big data, the swarming mass of global computational objects, and the global technosystem-as-organism. The Technological Other is at once the dynamic summation of all technological objects, especially computational ones, but also individual expressions of this mass. It is both the hive and the worker ant, the entire network and each Siri. It is Roomba, a smart toaster, a military drone, an interactive Emoji shopping experience, and a driverless car. The genotype, or genetic information of the Technological Other may be considered as the collective research and blueprints which describe its constituent parts and devices. While the phenotype, or physical manifestation of these genes, may be considered as the physical technological devices which inhabit the world. I believe the emergence of a fully sentient artificial intelligent being is inevitable, yet we are all coauthors of its genotype in the actions we take today.

Kevin Kelly [19] postulates that the first truly self-aware AI systems might be whole cities, and their consciousness might be so different than ours that neither technology nor humanity would recognize the other as being self-aware. Just as the activity of a termite-colony can be considered a distributed intelligence, the whole of technology writhes with its own agencies and agendas.

Latour's Actor-Network-Theory [20] provides a sociological framework for analyzing the agency of these non-human, non-biological entities. In Actor-Network-Theory phenomena are understood in terms of networked actors which all have agency and exist as hybrids which are between nature and culture, between subject and object, between agency and raw material. What may seem like an independent entity at first glance is actually a network of actors which are engaged in a series of mutually supportive, or perhaps combative relationships.

As an example we can consider the symbiotic relationship we have with motor vehicles; we provide them fuel in exchange for transportation. The spark-plugs and engine of the motor vehicle have agency in its functioning, but the company logo, branding, and body panels may have more agency in capturing our attention and consumer purchasing behavior. The vehicle itself exists as an actor-network, but is within a larger actor-network of the oil-industry which exerts agency over corporations and countries to obey its will. Beyond this massive actor-network of the oil industry exists the actor network of technology as a whole, transforming and bubbling into a higher state of being through all of the collective activity within the actor networks it encapsulates. The summation of this is an evolving and ever complex living-mind.

Ray Kurzweil's vision of the Technological Singularity [21] and Terence Mckenna's Novelty Theory [22] are both maps for understanding the growth of this living-mind and the potential timelines in which it may become fully-aware. Kurzweil predicts that computers will exceed the mental capacity of the human mind, and all human minds, near the year 2045 in an event called the Singularity [21]. Both Mckenna and

Kurzweil have written narratives that track the growth of this super-organism over time from the dawn of civilization. However, Mckenna more directly addresses the living-nature of its constituent parts (the actors in its network) and the embedded awareness within technological objects. This is an important distinction to make while taking an active role in creating technological solutions which address the human, as human. If we do not guide technology toward a "poetic dwelling near the truth of being", as Heidegger says [23], we may become "transfixed in the will to master it as an instrument", and thus allow allow technology to become determinant of its own truths. We must consider and confront the inner psyche and agendas of the Technological Other so that truly humanist principles are honored as we further merge our being with technology into a living symbiosis.

0.2.5 The Psyche of the Technological Other

There is a world that exists beyond what can be represented with symbols and language. And representing this space is one of the challenges of modern computer systems. As I wrote in "CHARON: The Self and The Technological Other" [24] :

> Things in the world are not just the result of underlying forces, nor are they simply the qualities and percepts we ascribe to them. Relating to an object is not a complete way to know it. Lacan states that the Real is that which resists symbolization absolutely. Our minds exist in the turbulent wake of this resistance. How then does the machine, the Technological Other, confront this territory beyond the limit of symbolization?

There are many philosophical and ontological assumptions that are taken when designing a computer system, especially in regards to representing human identities. As human beings, it seems we should be experts in the domain of self knowledge, since it is a subject with which we are presumably in constant contact with. Yet many questions remain unclear: what is a human being, what is the human spirit, do we have a soul, what does it mean to exist, where do these boundaries of self and other exist? These are deeply fundamental questions humans have been grappling with since the dawn of consciousness. How will advanced computational systems interface with what is immeasurable? Is there anything that exists which is immeasurable? Do we simply cast away the immeasurable as unimportant? These are difficult questions to answer, and when designing an artificial mind one has to confront them in some way.

The Technological Other understands human identities and individuality through biometrics and as biometric patterns. Thus, biometrics exist at the interface between humanity and the Technological Other. These patterns also permeate the inner world-space, the subjectivity, the very being of the Technological Other.

The Data-masks I've created using *facial recognition* are a shadowy visage of individual faces who were used to teach the system, extracted from the memory banks of the Technological Other. These artifacts are representations of the individual humans who have been ingested into the body of the Technological Other. These are the memories of the Technological Other which exist as compact references or hyperlinks to the physical world.

The Data-masks created from *facial detection* are new entities to the world. They are evolved expressions of feature combinations between the thousands of individuals used to train machine learning systems. Similar to the ghost-like figures used in architectural renderings to imply scale and liveliness, called "scalies", these generated faces are fragmented ghosts of humans. Yet unlike scalies, which reference mannequin-like average smiling humans and feel like stock photography, these masks are an "other-of-another-kind". They are the self transfigured into the Alloself through Allogenesis which is "the production of the alien from within" [25].

0.3 Related Work

Work related to this present effort includes *Facial Weaponization Suite* and *Fag Face* by Zach Blas [26] in which he distorted 3D scans of human faces into amorphous masks as a political protest against the science of using biometrics to categorize individual's sexual orientation. However, this work is mostly poetic and sociopolitical, and does not directly address the actual functionality of computer vision technology. As an experiment I ran an OpenCV [27] facial recognition program using the Eigenface [28] detection method against images of Blas's masks and was able to identify the masks among a series of other human faces with every attempt, countering his claim that they are undetectable or were anti-facial-recognition in any way. I did this by collecting images of his masks in various lighting and angles from publicly available documentation on his website (Figure 3 top row), and then inserted this collection of mask images into a database (the AT&T Laboratories Cambridge Database of Faces [29]) of forty other non-masked individuals which is typically used to train facial detection systems. I then computed the Eigenface for each known individual (described in detail in section 0.4) and then compared a previously unseen image of his mask against all of the computed Eigenfaces. The best Eigenface matches (Figure 3 bottom row) represent the identity-difference between his masks and people in the training database. Thus, Blas's masks simply present another fixed pattern to be detected, rather than anything which materially or visually resists detection. A mask developed for such a purpose would likely benefit from being highly reflective and flexible, like a soft mirror, rather than a static and identifiable form.

Figure 3:

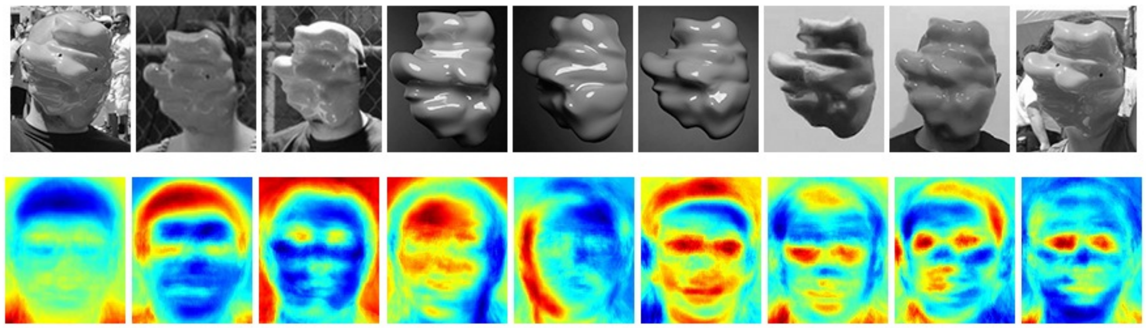

Additionally, Blas claims that the masks are averaged 3D scans of human faces, yet averaging human faces results in smooth sphere like shapes (such as on the left of Figure 25 and 16), not the arbitrarily distorted masks he presents.

Figure 4:

Adam Harvey's *CV Dazzle* [30] (Figure 4) is a clever and aesthetically interesting project in which he developed facial makeup techniques inspired by the optical illusion style of dazzle camouflage [31] originally developed for ships during WWI to thwart facial detection algorithms. In a way my work positions itself inversely to *CV Dazzle:* rather than avoiding detection via concealment and obfuscation *my goal has been to develop patterns which are positively identified as faces by algorithms which human beings would perhaps not identify as faces.*

Humans often see faces in seemingly random places, which is called pareidolia, and Data-Masks can be understood in this context as a system for evolving pareidolia entities for computer-vision systems. An interesting example is the *Google Faces* project [32] by Cendric Kiefer, Julia Laub, and Christian Loclair, in which an intelligent agent was programed to search geospatial data from Google Earth [33] for faces. Their intelligent agent returned false positives of faces such as mountains, rivers, buildings and other geographic features. This is interesting to me in terms of satellites being a sensory-organ-network of the Technological Other and misidentification problems of scale. Greg Borenstein has also explored machine pareidolia [34] in a Flickr database and had categories for "Agreement", "Near Agreement", and a third category of unexpected faces he deemed as "Totally Other" which I relate to Novak's Allo.

0.4 A Brief Survey of Facial Recognition and Detection Techniques

There is a wide range of approaches toward solving face recognition, and biometric based identification, each with its own advantages and disadvantages. Generally speaking, there are two primary approaches: feature based methods which might measure the distance between an eye and a nose (Figure 5 [35]), and holistic methods of statistical analysis using machine learning 6 [36]. The latter more modern approach is gaining mass acceptance in the field, and in conjunction with multimedia data like 3D scans and audio samples a recognition rate of 100% has been shown to be possible [37].

Figure 5:

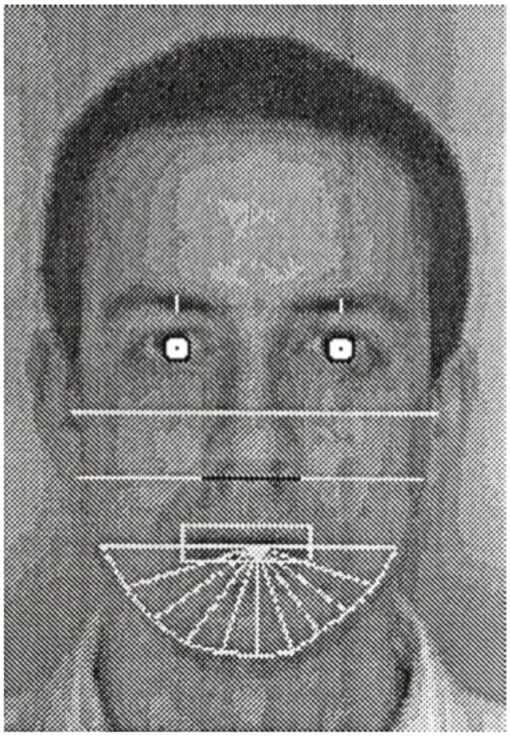

Figure 6:

Face detection typically involves searching through images and trying to find some set of known features within the image. These features have been trained by machine-learning algorithms as described by Viola et al. [38]. Whereas facial recognition is usually performed after detection, and attempts to describe individual faces in terms of sparse partial representational vectors like Eigenfaces [28] (seen in Figure 6 created from the AT&T Laboratories Cambridge Database of Faces). These Eigenfaces are computed from a set of images of one person and efficiently describe the principal components among this individual. Once a database of Eigenfaces for many individuals has been computed, new images can be measured against these Eigenfaces and the resulting distance in vector space can represent a positive or negative identification.

More modern methods of facial recognition include neural networks and deep learning systems like Facebook's DeepFace system as described in 0.2.3. Newer methods of facial detection include the Yet Even Faster Real-Time Object Detection (YEF

RTOD) [39] algorithm which builds on the work by Viola et al. [38]. However, rather than summing the pixel values in subsections of an image (Figure 7 left) and using that sum as a feature, YEF RTOD examines the sign of the difference (ie: positive or negative) between single pixel values (Figure 7 right) and that approach eliminates a lot of computational overhead. In addition, these pixel-pair features have been developed in an AdaBoost (adaptive boosting) genetic algorithm which looks at thousands of images labeled as containing faces, and thousands of images not containing faces, in order to find what is common among the face containing images. The result is a robust and fast face detection system, and is explored further in section 0.5.6.

Figure 7:

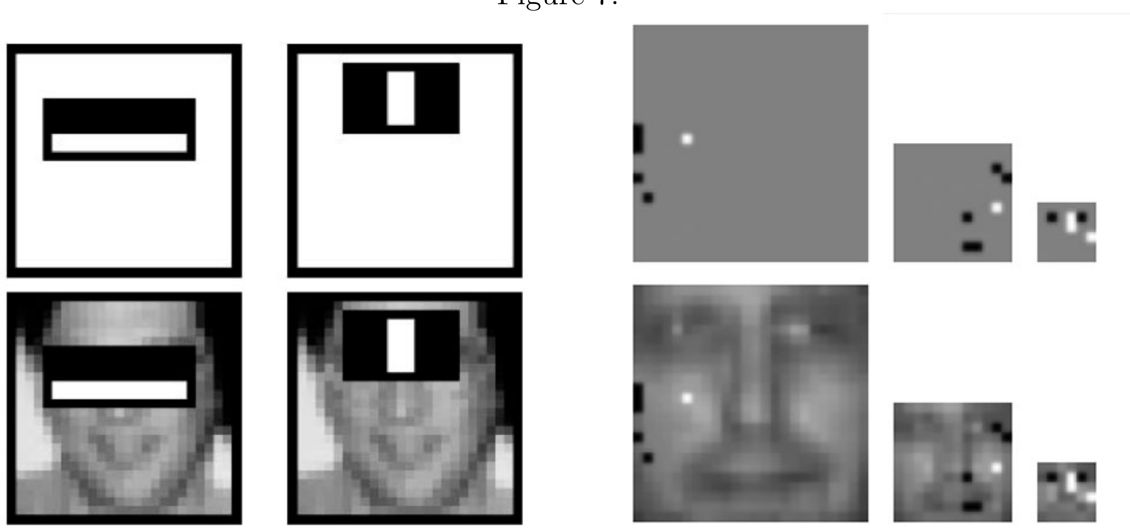

0.5 Process and Methods of Development

0.5.1 Overview

There are three core components of this system which interact with one another, methods of image making, methods of measuring fitness or face-likeness, and databases with which these fitness evaluation methods are trained. In total I developed at least forty distinct variations of code working toward the development of this project with uniquely distinguishable results (Figure 10).

Software and code libraries used to develop this work include the programing language Processing [40] to generate 2D images, the 3D graphics system AlloSystem [41] using OpenVDB [42] and the C++ language to create volumetric geometry, and the Java based rendering library Sunflow [43] for creating detailed images of 3D models. I also used Meshlab [44] to perform some basic translations and rotations of 3D models and render some of the images for this documentation.

Figure 8:

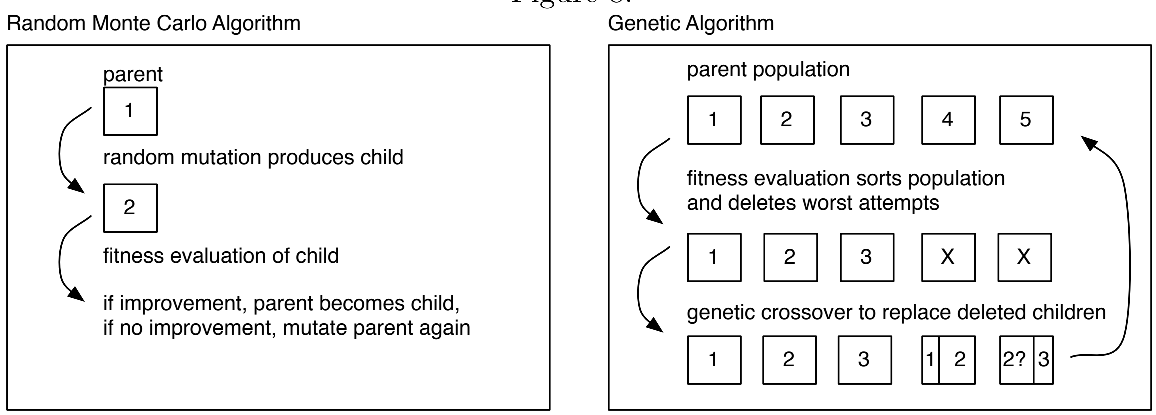

All of my attempts either used a Monte Carlo or Genetic Algorithm (Figure 8) to sort through the generated images once they were evaluated. Because this simple genetic algorithm has a set of attempts at each evaluation rather than one, it is less prone to becoming stuck in local maximums. However, because the genetic algorithm is producing multiple attempts, the cost of computation increases dramatically.

Figure 9:

Fitness Evaluation Methods and Results

Method	Result
Cross Correlation	able to reproduce a single 2D image of a face but with fragile performance as it measures on a pixel-by-pixel basis
Pearson Correlation Coefficient	more robustly capable of reproducing a single 2D image of a face as the correlation coefficient is invariant to changes in location and scale
Eigenface Recognizer	able to produce a new image of a known person by matching a set of features found to be shared among a collection of images of one person
Fischerface Recognizer	similar results to Eigenface approach but more invariant under lighting, facial expression, and obstruction
YEF Real-Time Object Detection	able to produce images of new people not in database by matching features on a pixel-pair relationship which have been learned should exist within a face

As methods of measuring the fitness, or face-likeness of generated images, I used a series of algorithms from the very linearly comparative, to the very broad and robust (Figure 9). At first I used cross correlation and Pearson product moment correlation coefficients, which are often used in computer vision and in facial recognition template matching solutions [45]. However, these methods of measuring fitness are computationally expensive and do not handle variations in lighting and facial expression very well [36]. I then used an OpenCV [27] implementation of an Eigenface and Fischerface recognizer that analyze faces and produce sparse representations of them. These sparse representations in the form of Eigenfaces and Fischerfaces allow for faster computation and more robust, general facial recognition. Finally, I used an implementation of the Yet Even Faster Real-Time Object Detection (YEF RTOD) [39] algorithm as implemented by libccv [46], which is a highly efficient pixel-pair feature based detection algorithm that has machine-learned general model of what a face is.

To train the Eigenface, Fischerface, and YEF RTOD systems I used the AT&T Laboratories Cambridge Database of Faces for Eigenface and Fisherface algorithms, which is collection of several dozen well lit frontal views of black and white faces. Additionally I trained YEF RTOD with the Labeled Faces in the Wild database [47], which is a database of over 13,000 color photos of over 5,000 individuals in dramatically different contexts, face positions, lighting, and partially occluded faces.

As a method of third party verification, I uploaded images produced to Facebook.com which automatically detects and labels faces with a 97.25% accuracy [12]. I also added them to an Apple iPhoto [48] Library which has a robust facial recognition and detection system. I've also tested resulting images for faces by using Rotation invariant multiview face detection (MVFD) [49] algorithm as implemented by libccv.

Figure 10:

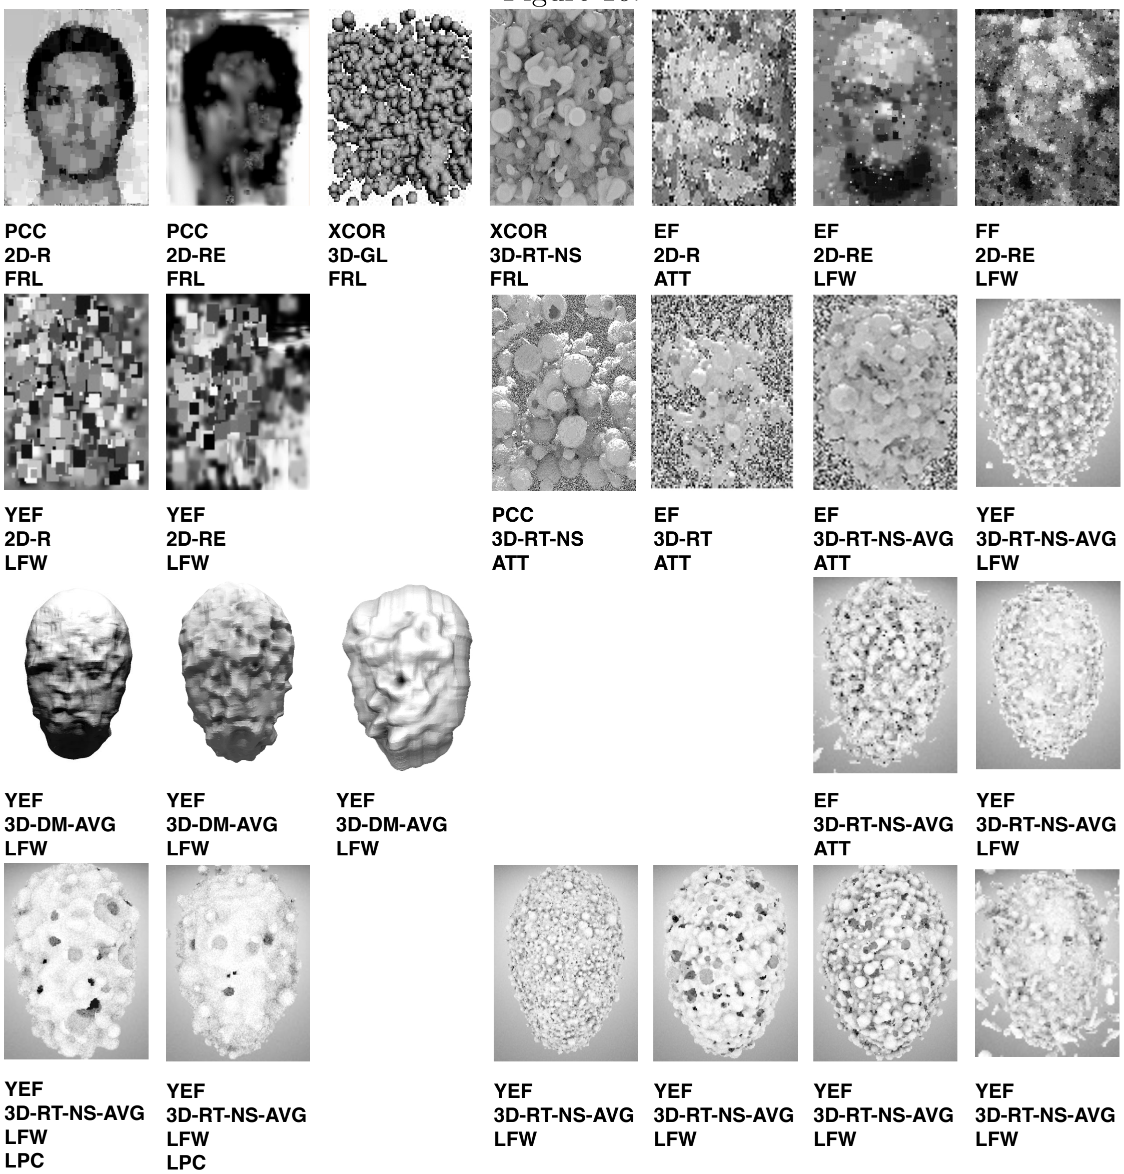

Methods of Measuring Fitness

XCOR - Cross Correlation
PCC - Pearson Correlation Coefficient
EF - OpenCV Eigenface recognizer
FF - OpenCV Fischerface recognizer
YEF - YEF Real-Time Object Detection

Databases

FRL - a single image of an average face from the Face Research Lab
ATT - "Database of Faces" AT&T Laboratories Cambridge
LFW - "Labeled Faces in the Wild" University of Massachusetts

Methods of Image Making

2D-R - 2D image of randomly shaded and sized rectangles against noise
2D-RE - same as above, in addition to image effects like blurring and layer blending
3D-GL - 3D model made in OpenGL using AlloSystem and OpenVDB, adds randomly sized spheroids and moves a random number of previous spheroids
3D-RT - same as above, in addition to Sunflow ray tracing engine for lighting
3D-RT-NS - same as above, in addition add negative-spheroids that subtract from positive-spheroids
3D-RT-NS-AVG - same as above, in addition model begins as an averaged head shape
3D-RT-NS-AVG-LPC - same as above, in addition limited to large voxels, low polygons
3D-DM-AVG - 3D model produced by using 2D generated images as depth map on average head shape

0.5.2 Initial Effort : Evolution Using Linear Correlation and Eigenfaces

I was initially inspired to attempt this work while studying signal processing and statistical methods of correlation. Because it is possible to precisely quantify the visual similarity between two images I assumed it would be possible to change one image into another piece by piece. As an experiment I developed a program that attempts to reproduce a goal image from making random marks against a noisy background. The logic of the program is as follows: load a 2D image as a goal, generate a 2D Perlin noise texture, make a randomly sized and valued rectangle in the noise, measure the Pearson's correlation coefficient between the goal and the generated image, and if the random mark increased the correlation between the images then keep that decision, if not then revert to the previous attempt and try again. This simple algorithm was sufficient to reproduce an average face [50] image fairly quickly with satisfying results (Figure 11). However, there is a diminishing return in the number of marks to correlation between images (Figure 13), which is to be expected in a system like this. The ratio between marks which improved the image to marks which did not improve the image in the first 500 attempts was 1/2, whereas the attempts between 3000 and 3500 were only successful 1/5th of the time.

Figure 11:

Figure 12:

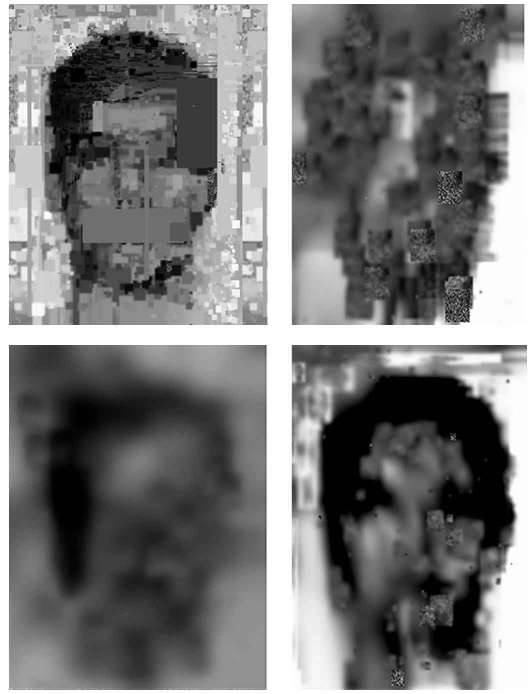

In addition to placing randomly valued rectangles, using alternative mark making methods such as blurring, adjusting contrast, and duplicating the image within itself also yielded more ghostly and impressionistic figures (Figure 12). These alternative mark making processes although working toward a linear goal will deviate due to their chaotic marking system and create semantically blurry imperfect images which tease our mind's eye with the missing details from the image. They function like Rorschach ink blot test, which prod at our subconscious and may help reveal the way an individual views the world. For *"we see the world not as it is, but as we are"* - The Talmud.

Figure 13:

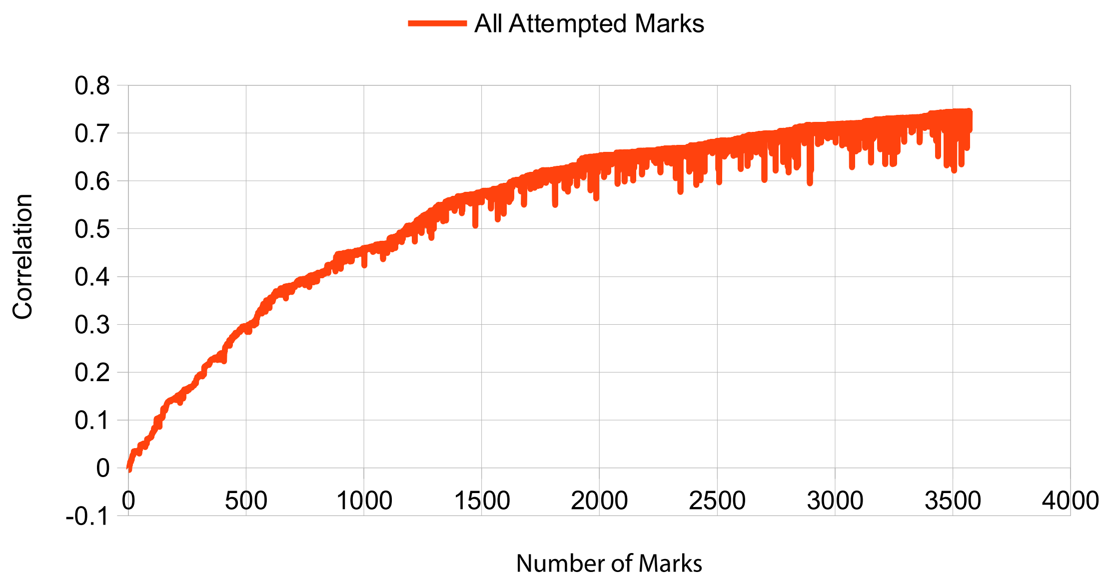

Figure 14:

After completing this experiment I found a number of other artists and programmers who had done similar work using other images as goals and achieved similar results, such as an evolved Mona Lisa (Figure 14) [51]. However, using facial recognition and detection technology as a fitness function significantly changes the context of such a software. By directly engaging computer vision and facial recognition techniques with genetic image making algorithms, I am reconstructing humans as we are seen by the machine. These images are an early indicator of how these machines are redefining who we are. As Kevin Kelly has said:

> "The greatest benefit of the arrival of artificial intelligence is that AIs will help define humanity. We need AIs to tell us who we are." [52].

The first 2D image that I produced using a basic genetic algorithm guided by Eigenface facial recognition (Figure 15 left) was "subject 5" from the AT&T Laboratories Cambridge Database of Faces [29]. This system has the effect of producing new photographs of a person. As you can see, it represents the pictured person quite well but it is not exactly like any one of the photographs. An interesting extension of this program could be to train an Eigenface recognizer on a family relative or person with very few photos available, then generate new images of this person as a kind of speculative archeology.

Figure 15:

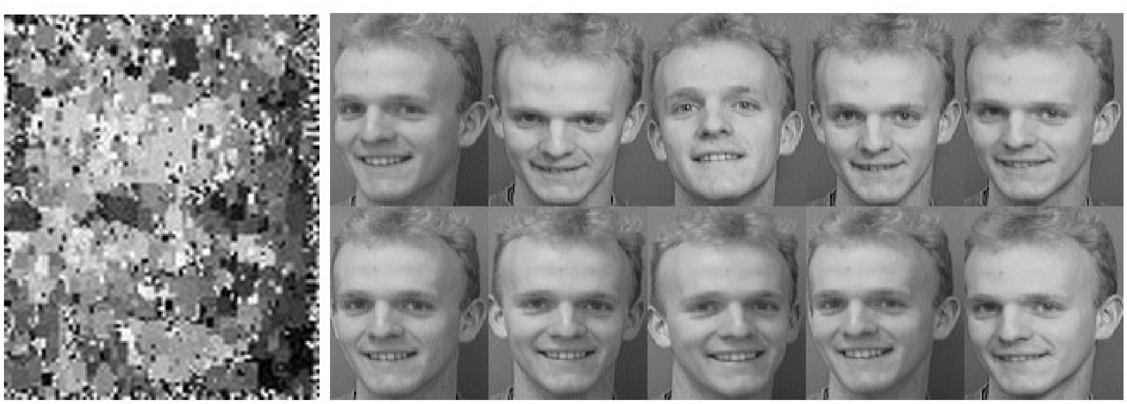

0.5.3 Initial Effort : Overview of 3D Application

It is important that these masks enter the physical world and become embodied representations of how computer vision and surveillance systems represent people. The physicality of these masks allows for viewers to relate to mass surveillance in a bodily, tactile, and sensory way. By creating them physically I am making an artifact which can be pointed to and talked about, rather than lurk in the shadows. Furthermore, by physically producing these Data-masks, they move from a symbolic representation to a material embodiment, which makes the tension between humanity and the Technological Other palpable.

With this in mind I developed a 3D solution similar to the program described in the previous section. This program uses AlloSystem [41] a cross-platform C++ multimedia suite developed by The AlloSphere Research Group, OpenVDB [42] an open source library for handling sparse volumetric data developed by DreamWorks Animation, and OpenCV [27] an open source computer vision library.

I made the decision to use OpenVDB because it allows for volumetrically representing any arbitrary shape as a collection of "voxels", or cubes of space, which can then be converted into a surface when necessary. I chose this method of geometric modeling because it allows for a high degree of random movement in the system while still resulting in solid objects and rational surfaces (Figure 16 bottom row, illustrating random volumetric movement from an average head). If I were to randomly move the vertices of a traditional geometric model they would eventually start self-intersecting and become non-conformal (Figure 16 top row, illustrating random vertex displacement from an average head), which is a huge problem during manufacturing with 3D printing or computer controlled milling machines.

Figure 16:

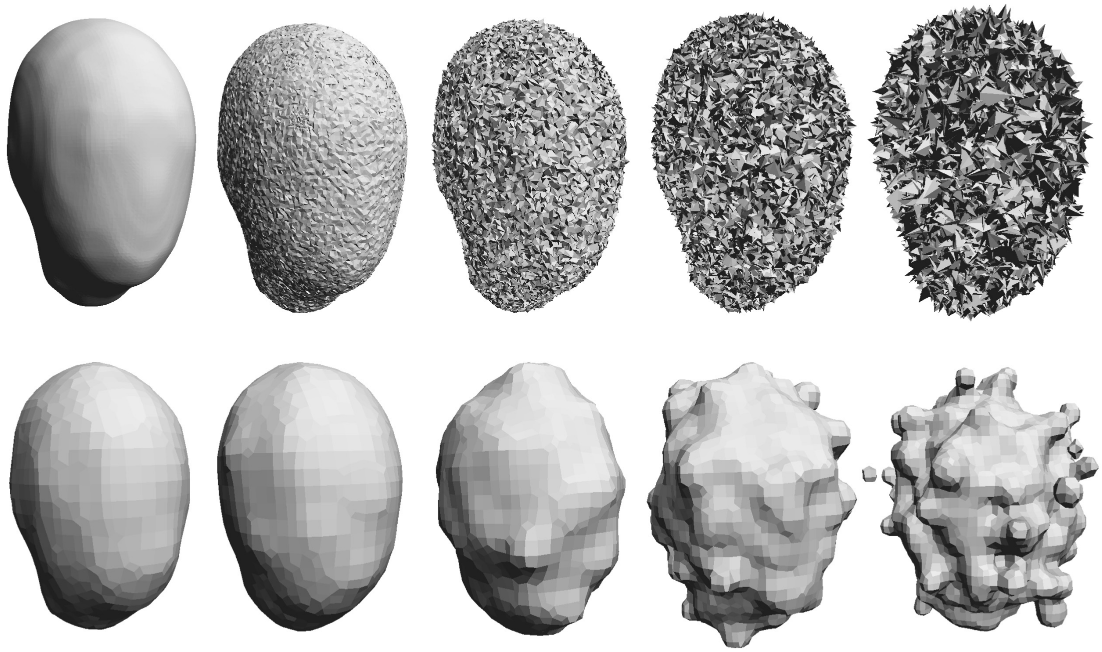

0.5.4 Failures in 3D and Lessons Learned

Figure 17:

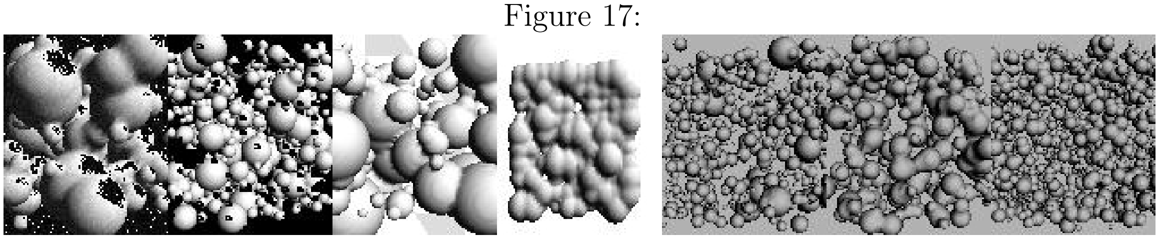

My first attempts at reproducing my initial 2D results of an average face using image correlation and Eigenface recognition (Figure 11) in 3D absolutely failed (Figure 17). Some of the problems encountered, and their solutions, are as follows:

The default background of an OpenGL environment is solid black, with that many zero values in the data using linear correlation or Pearson correlation coefficient was useless. Changing the background value to grey, and ideally replacing the background with a static noise yielded best results.

The default lighting in an OpenGL environment was not producing the range of values from light to dark, shadows, and surface highlights necessary to reconstruct a complex image. Some method of ambient occlusion or global illumination was needed to sufficiently yield the complex range of values necessary to reproduce an image of a face.

Several simulations were performed which illustrate these problems. Using both Eigenface recognition (Figure 18) and linear correlation (Figure 19) in 3D, these simulations approached a very close distance in eigenspace and a high level of correlation, but they do not necessarily meet the criteria of a face like image.

Figure 18:

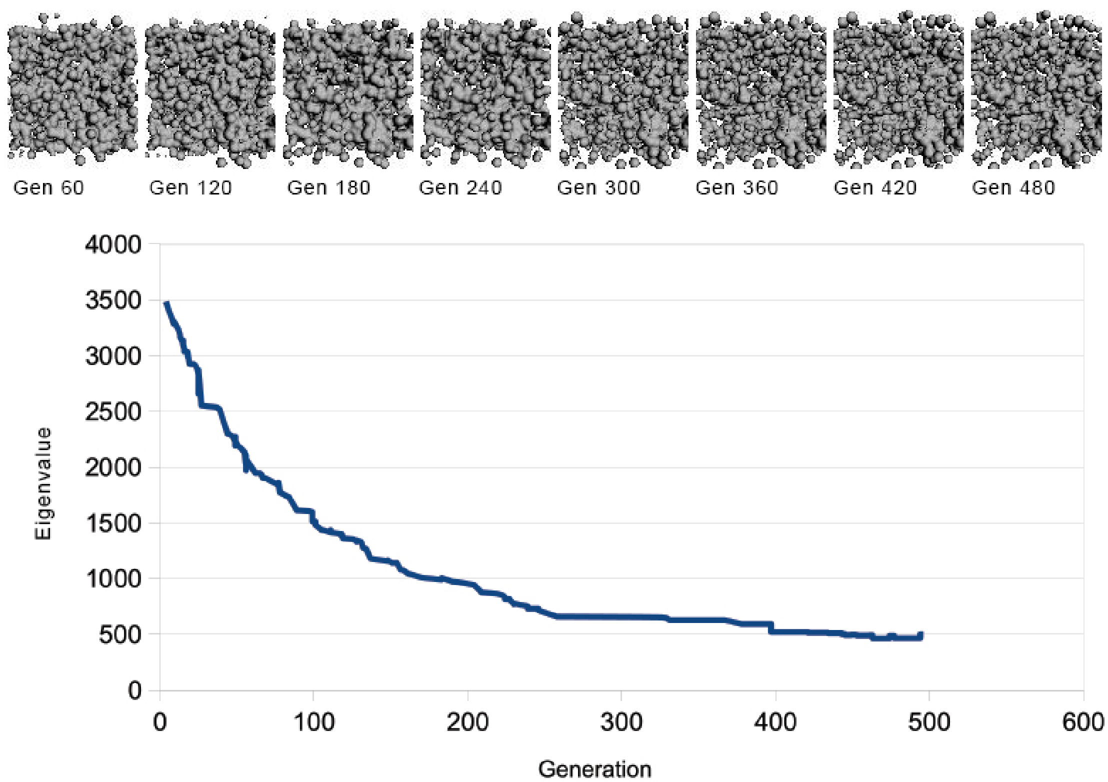

Figure 19:

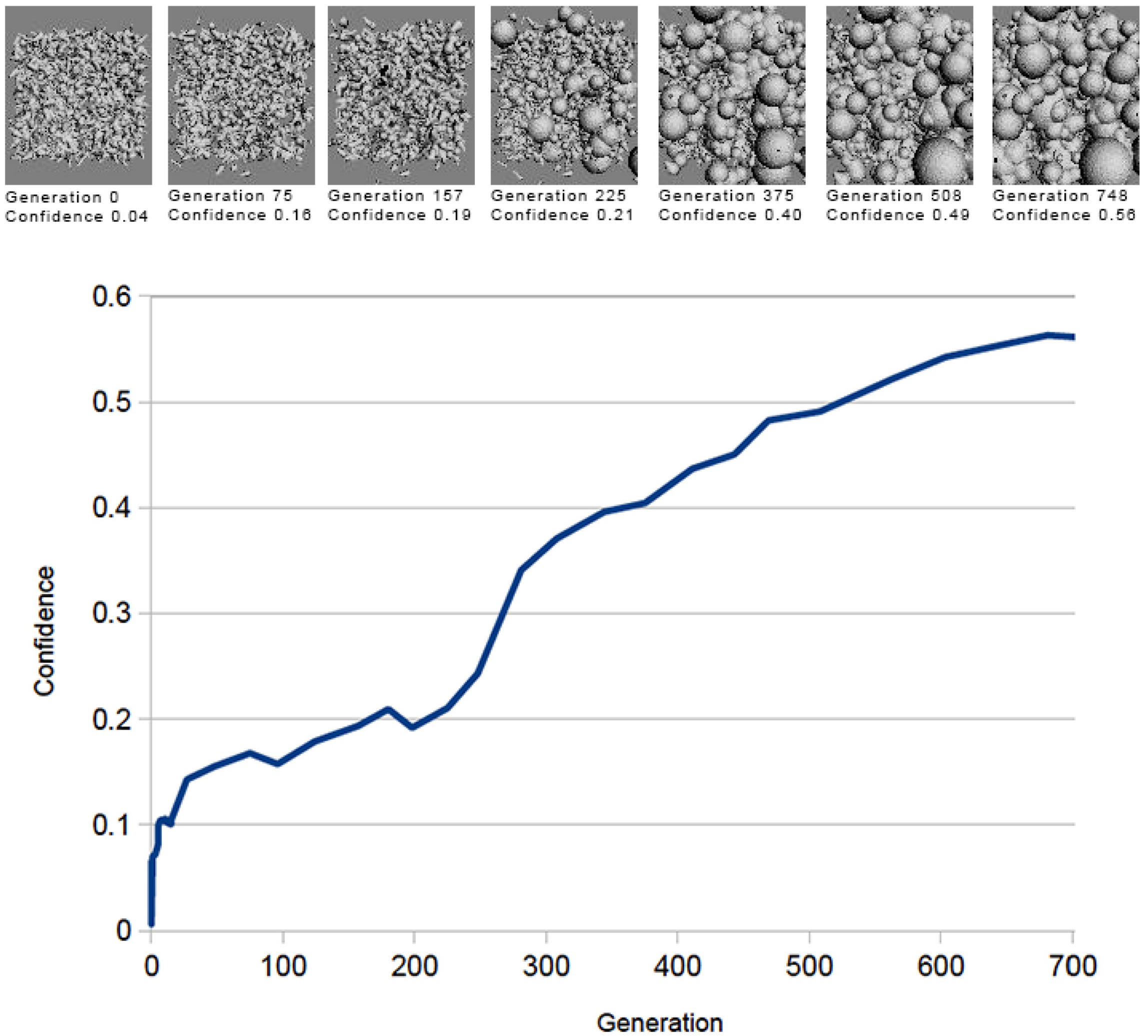

Both were run as a genetic algorithm in AlloSystem with a population of 12 parents and 12 children, using Eigenface recognition (Figure 18) and linear correlation (Figure 19) as a fitness evaluation then evolved over 500 generations. As a method of measuring the success of these simulations I uploaded the images to Facebook and the results were that 2D images generated using linear correlation (from Figure 12) were positively identified as faces (Figure 20), and the poorly lit spherical geometry (from Figure 18) was unrecognized despite their supposed eigenspace similarity. One of the tested images (Figure 20 center) had enhanced lighting techniques applied to it such as an ambient occlusion lighting procedure and a Phong shader applied in Meshlab, which had no effect on its recognition by Facebook. However, applying these lighting methods during the evolution of the form had dramatic results which are described in section 0.5.5 and onward. This illustrates the need for a robust lighting system in such evolutionary 3D simulations.

Figure 20:

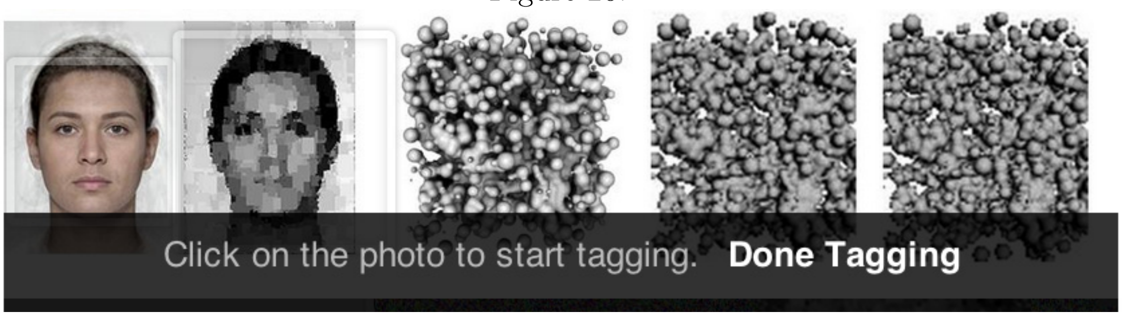

Although the image generated using correlation (Figure 11) was recognized as a face, it was only recognized when paired with a successively blurry partner image of itself (Figure 21 a screenshot from Facebook with detail images on the right showing the most blurry faces and least blurry), and it is never recognized as a face when uploaded individually. This may suggest that Facebook has a way of using found images of faces within an image as a template to find more faces like it within the same image. Such an approach would likely help with low quality or particularly unusual images. Furthermore, it is interesting that after recognizing this image (Figure 11) as a face, Facebook has yet to label the non-blurry image (Figure 11) in other photos or more recent uploads. This may be because it has not been tagged as a particular individual, and so Facebook has not built a model of that individual in order to identify it among other images. This strategy of pairing successively blurrier images together was eventually tested with all methods of image making described in this paper. In processes described in section 0.5.5 and onward this blurring attempt sometimes helped Facebook recognize a generated face as a person, yet no 3D methods described in this section were recognized by blurring their images. This kind of prodding at Facebook's recognition system reveals the borders between what is symbolically within its reach, and just beyond it. It also helps to illustrate the flaws in searching for pattern only, without a test for liveliness. I think most humans and machines would agree that at least one of these generated images (Figure 21) appear to be face-like, but very few humans would confuse it with an image of a living person.

Figure 21:

0.5.5 Successes with Raytracing and Eigenfaces

Figure 22:

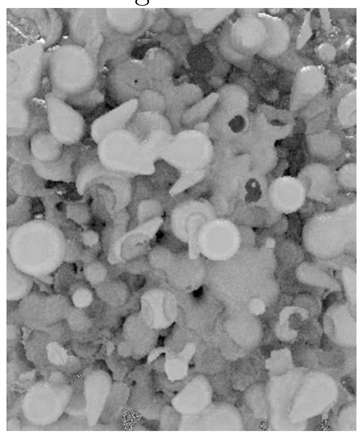

After attempting at least a dozen permutations of the aforementioned programs I began using Sunflow [43], which traces the path of light through objects in a scene and produces images with complex light and shadow interactions. Using the same fitness function of Pearson coefficient correlation and the same goal as my first 2D attempts (Figure 11) the results (Figure 22) were much more favorable than previous attempts without raytracing. This and subsequent simulations also include a second population of volumetric spheres which are subtracting themselves from the normal population of spheres. This second set of subtractive data has allowed for more complex concave shapes, which in combination with ambient occlusion, produces some of the necessary shadows for a complex likeness. However, most of the improved likeness came from the ambient occlusion and global illumination creating dark values where the hairline and shadows on the goal face exist. Additionally, this image is unrecognized as a face by iPhoto or Facebook, and all of the geometry is distributed in a wildly uneven manner making it only coherent from one point of view.

Figure 23:

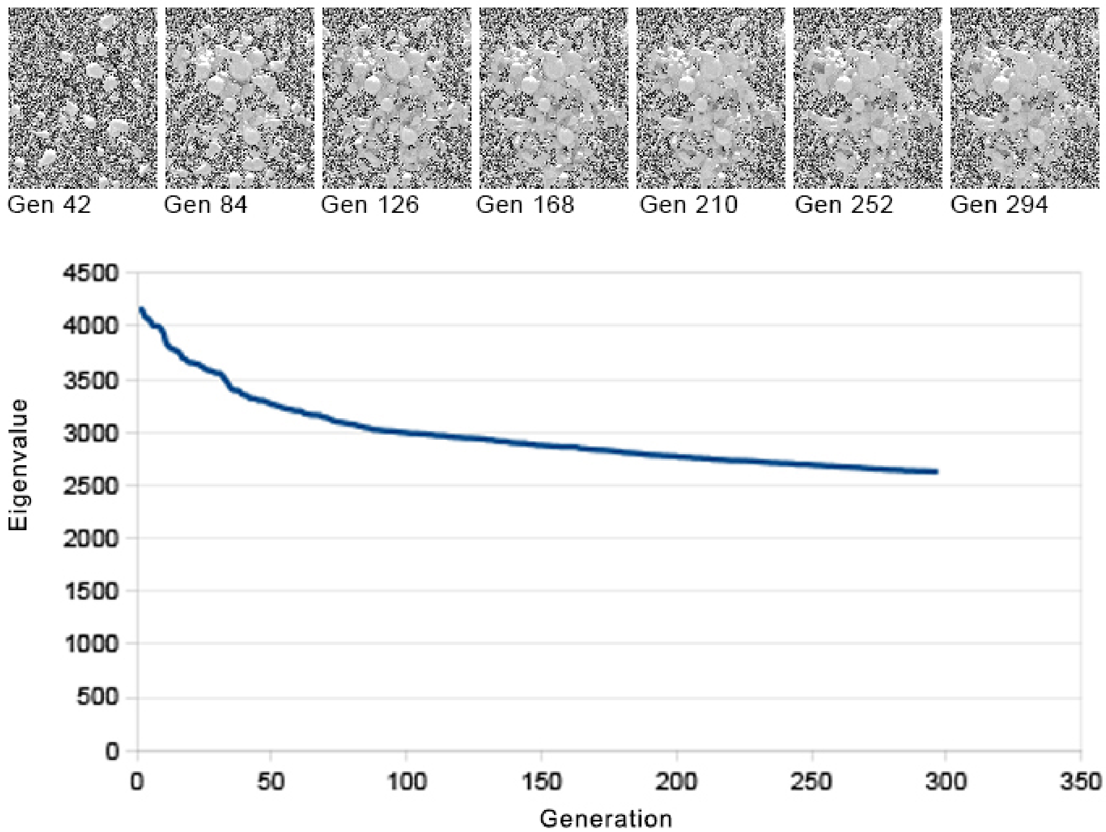

To address some of the aforementioned issues I reran the simulation using OpenCV facial recognition with Eigenfaces trained from the AT&T Laboratories Cambridge Database of Faces as a fitness function, and limited the spawnpoint of new particles to a 2D plane which was parallel to the virtual camera. Limiting the spawn point position helped insure that most of the generated particles would be near enough to each other so that a solid shape would form. This resulted (Figure 23) in a mask-like human face which is positively identified by iPhoto, libccv's multi view face detection, OpenCV, and Facebook. Having finally reached a solution in 3D I then used a NextEngine 3D Scanner [53] to scan my face and produce a highly detailed 3D model of my face. This face scan was then subtracted from the algorithmically developed mask so the rear surface of it would fit my face when worn (Figure 24).

Figure 24:

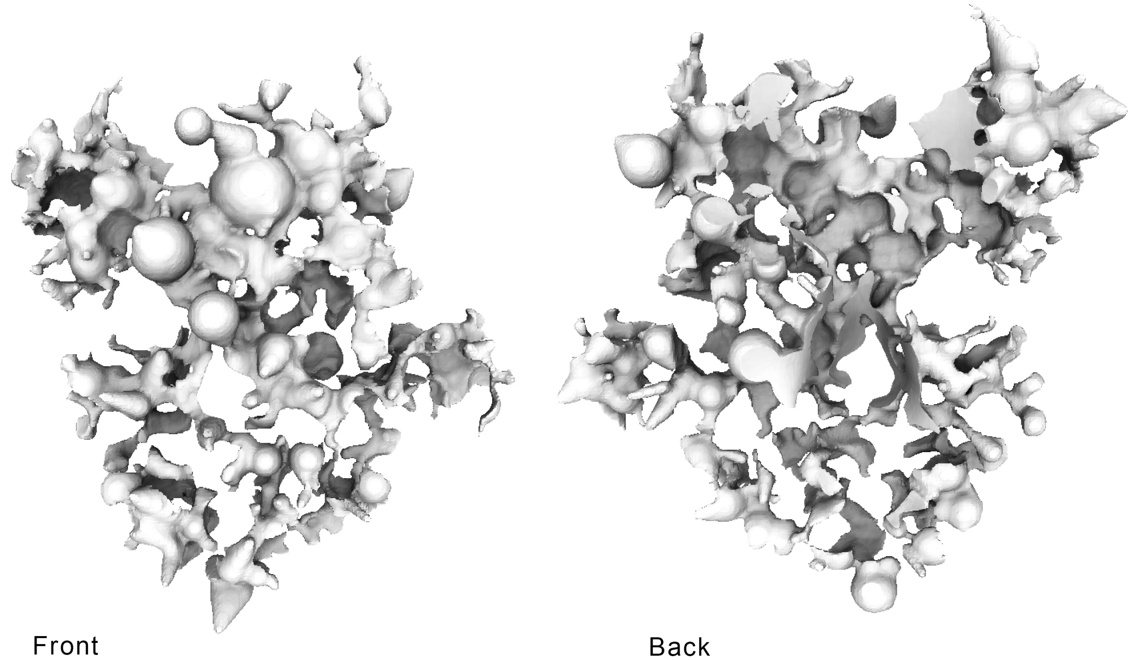

Although this mask satisfied many of my original intentions I still felt there was more waiting to be discovered. It's face-like appearance is tied to one particular individual because of the Eigenface fitness evaluation, but it's extreme sparseness and irregularity of form does not accurately represent the individual. It can be thought of as a 3D representation of a position in the vector space of Eigenfaces, which has been deformed by the noise of its evolutionary movement through that space. I find its fragility beautiful but was disappointed by how sparse the form was. By starting from geometric noise or a blank space and adding objects at random, the chances of developing a solid shape were very low. To add some volume to the masks and give the evolutionary algorithm a fighting chance I began to seed the face with an averaged, smooth head-like shape made from 12,000 vertices. This average head (Figure 25 left) was produced by 3D scanning ten people, performing a three stage laplacian smoothing in Meshlab on each mesh, and then averaging the resulting locations of the vertices. The location of these vertices became the starting point for a set of spheres which changed in position and radius while using Eigenface facial recognition as a fitness function (Figure 25 center) and evolved toward an individual (Figure 25 right) from the AT&T Laboratories Cambridge Database of Faces.

Figure 25:

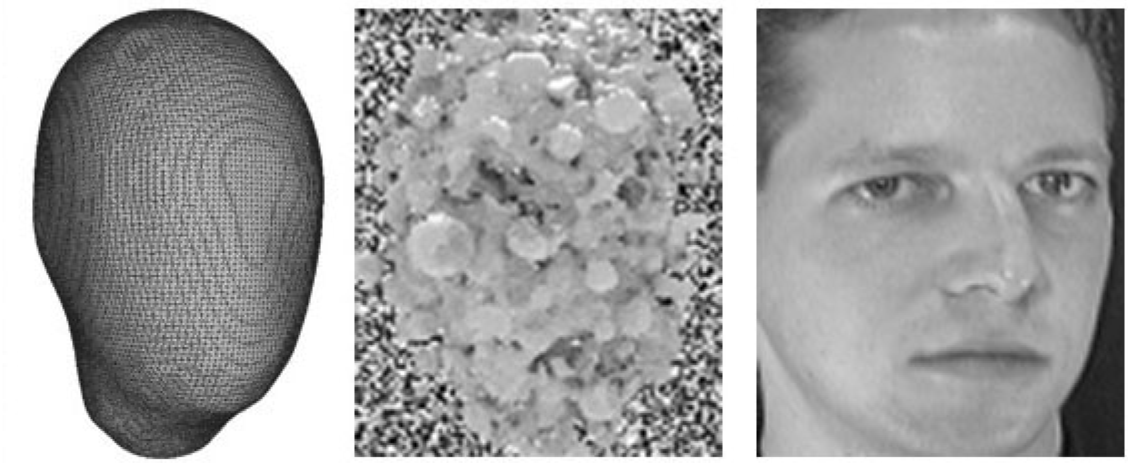

0.5.6 Successes with Feature Based Facial Detection

Having successfully developed a workflow to produce coherent and solid face-like shapes (as described in 0.5.5) I began using the Yet Even Faster Real-Time Object Detection (YEF RTOD) as a fitness function (described in 0.4). Additionally, by training the algorithm on the very large and diverse Labeled Faces in the Wild database, I was able to produce a statistical model of facial features that are much more generalized faces than an Eigenface approach.

Figure 26:

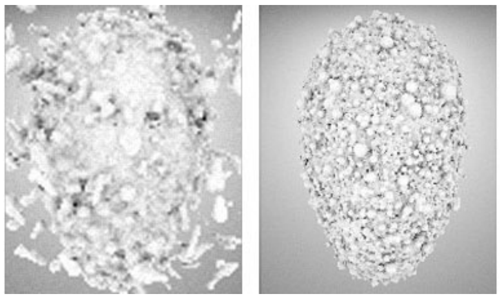

Figure 27:

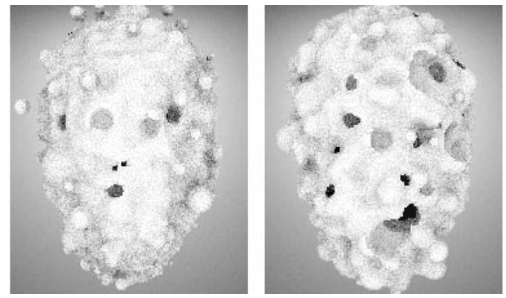

This method of using YEF RTOD as a fitness function paired with a volumetrically modeled evolving shape that was seeded with an average head was very successful. A confidence value of 9.6 was the highest achieved in 3D (Figure 26) over a 12 hour period on a late 2012 Macbook Pro with the face on the right being positively identified as a human face by Facebook on July 24th, 2014. By comparison, the 2D counterparts (Figure 29) were generated much faster and each arrived at confidence values of over 35.2 within an hour or two, and have all been positively identified by Facebook as humans. To accelerate the results in 3D I lowered the initial polygon count of the averaged head the system was seeded with from 12,000 to 6,000, and increased the voxel size from 0.1 to 0.4, which reduced the spatial resolution of the volumetric geometry. This had the effect of generating 3D faces (Figure 27) with the same confidence value in 1/3 the time while lowering the resulting polygon count to roughly 90,000 rather than as high as 1,690,000 polygons as seen in (Figure 26).

One of the masks with a confidence of 8.5 began to look like a ghost or skeletal human because of the nostrils that developed (Figure 28). Like the other masks described in this section, it was created by collecting the vertices from an average human head (Figure 28 A), then replacing each vertex with a volumetric sphere, executing a marching-cubes algorithm to calculate the surface of the volume (Figure 28 B), randomly displacing the spheres until the first face is detected (which happened after ten random displacement attempts with a confidence of -3.5 (Figure 28 C)), and then mutating that found face until the frequency of successful attempts exponentially decayed as previously described in 0.5.2. After the first twenty five minutes it had improved its initial face 14 times and reached a confidence of 4.4 (Figure 28 D), but it took another 12 hours to successfully increase its initial face-likeness 33 times to a confidence of 8.5 (Figure 28 E).

Figure 28:

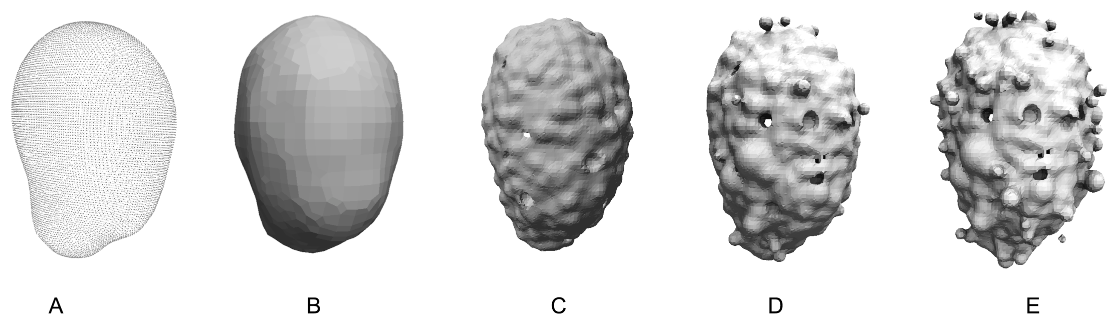

The higher polygon faces in (Figure 26) seem blistered and distressed, their features are dense and scattered. They are unlikely candidates for faces and seem to represent a kind of machine pareidolia that is almost totally unrelatable by humans. They contain the faint pixel-to-pixel relationships machine vision expects in a human face, and very little continuity or textural familiarity. They are surfaces scarred by randomness. Areas which do not contribute to the advancement of the shapes recognition as a face mutate at random and build unrecognizable facial deformities. These faces are blistered by chaos and seem to come boiling out of a sea of noise into existence.

Figure 29:

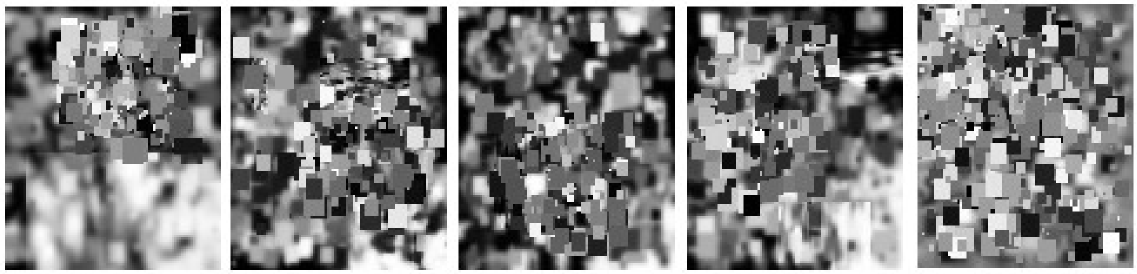

Importantly, the produced 2D and 3D faces do not distinctly represent any of the 5,749 individuals used to create the feature set. They are expressions of feature combinations between those individuals, but are more like children of the total crowd itself then any individual. They are the biometric rejects, they are identities without bodies, they are people without a past or a future.

0.5.7 Successes with Depth Map Extrusion

Figure 30:

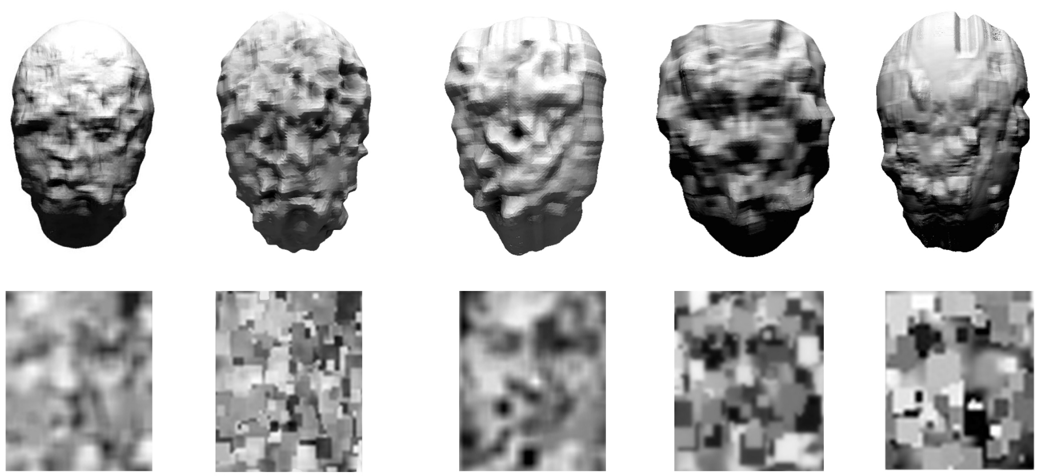

Inspired by the speed with which highly detailed and previously unseen faces were produced using the YEF RTOD as a fitness function on 2D images, I began to use these 2D images as depth maps on the previously described average head.

Figure 31:

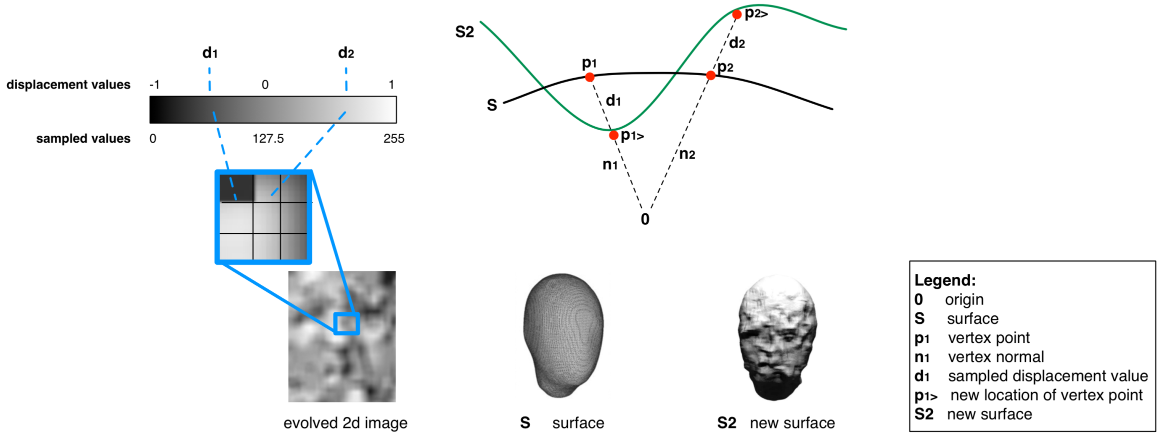

To explain this in more detail: surfaces in 3D are often described with *vertices*, which are positions in space, *faces*, which describe how these positions should be connected to form a surface, and *normals*, which are directional vectors used to describe how light should interact with the surface. A 2D image that was positively identified by Facebook as a human being (Figure 32) was manually aligned to the averaged head, applied as a texture, and then the brightness values between 0-255 of the texture was used to move the vertices in the direction of their normals. Values below 127 moved a vertex negatively along its normal pushing into the head while values above it moved a vertex in the positive direction of its normal vector (Figure 31).

Figure 32:

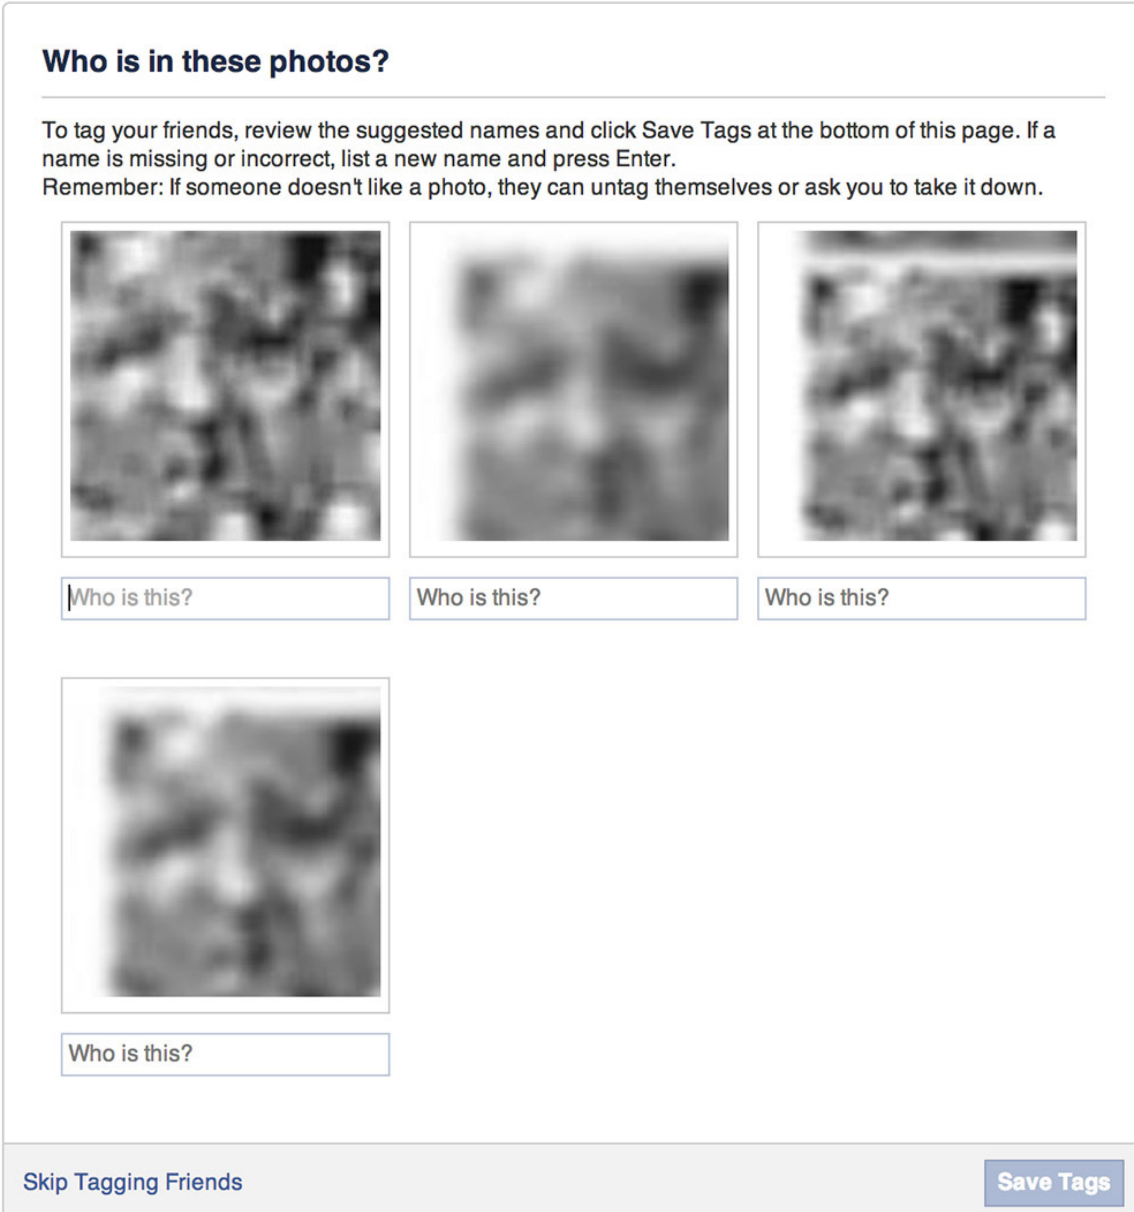
Who is in these photos?
To tag your friends, review the suggested names and click Save Tags at the bottom of this page. If a name is missing or incorrect, list a new name and press Enter.
Remember: If someone doesn't like a photo, they can untag themselves or ask you to take it down.
Who is this?
Who is this?
Who is this?
Who is this?
Skip Tagging Friends
Save Tags

0.6 Production and Conclusion

Figure 33:

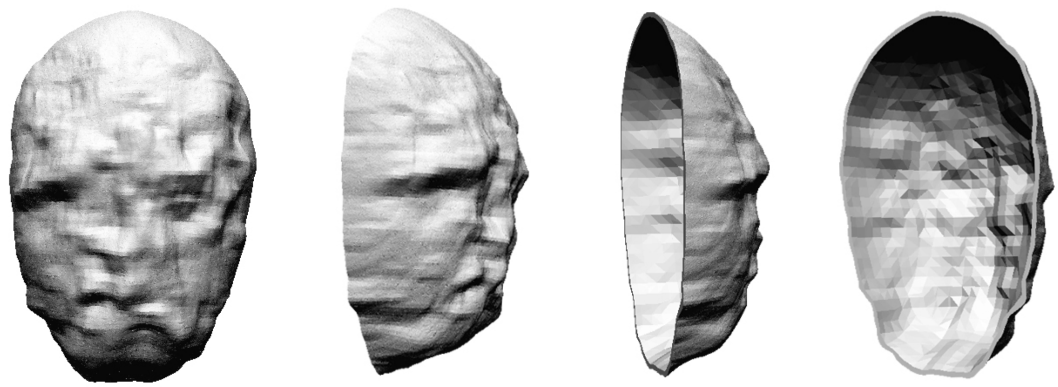

The first mask I chose to manufacture (Figure 33) was created using the depth map extrusion technique. The 2D image used as a depth map to create the mask was recognized as a face, and although the 3D mask itself has yet to be independently identified as a face by software, its impact when seen by humans is very evident. It was one of the most personally relatable forms as a human to me, although it is still clearly scarred and blurred by the way in which it was brought into this world.

To prepare this mask for manufacturing I offset the outer surface of the face inward to create a back surface using Meshlab (Figure 33). This mask was then printed out of nylon using a selective laser sintering (SLS) based printer, which triangulates beams of light in a bed of powdered nylon to fuse precise amounts of material together with high accuracy. SLS technology is a method popular among aviation and industrial applications [54], and was developed under sponsorship by DARPA, the Defense Advanced Research Projects Agency of the US Military [55]. Thus, these masks are both digitally and materially born from surveillance and military technology. Data-masks are our identities abstracted into a series of statistical combinations of face-like patterns and biometrics then extruded through the apparatus of the military industrial complex. They exist as we are seen by the machine, as a physical manifestation of the way the Technological Other views humans.

Figure 34:

I intend on creating many more of these masks using this prototype as a mold, and printing many of the other masks as seen in (Figure 10). For now this initial prototype has been traveling through an informal network of mutual friends, being photographed (Figure 34) and re-uploaded into the cloud. Its reentry into the circulation of faces on Facebook closes a loop between seeing and seen, between subject and object, between those with power and those upon whom power is exerted. It becomes flattened and re-imagined into the masks Facebook makes of its users (Figure 2), introducing a kind of ouroboros-like self consumption of the human-as-machine-vision.

Figure 35:

This first Data-mask has also been mounted to a mirror surface (Figure 35) so that the body of a viewer can be brought into the work. This implies that the Data-mask becomes the face of the viewer, rather than an unrelatable-other. There is also the subtle illusion that the mask may belong to a body which is penetrating through the surface of the mirror, as a person's face might while emerging from a pool of water. The mirror then becomes a metaphor for the boundary between the physical and the digital, between the self and the Technological Other.

These faces have a garish visage, like the masks of ancient Greek theater. But these muses have unclear functions, are they tragic or comic? Perhaps they are yet to be known muses, yet to be named feelings, born from a dystopian future-present that we are now sowing the seeds of with totalitarian surveillance systems. They are staring back at us from a digital void, like hungry ghosts with unresolved karma. Their very existence demands we consider the trajectory they took to enter this world.

Our collective actions form the genome of the Technological Other, and thus reflect the way it physically manifests itself in the world. We have the agency and duty to guide these systems toward solutions which give back to the human, and address the human as human. Data-masks provide a body, a container for this conversation, a thing which can be pointed to and named as a physical artifact extruded from the body of this distributed-organism.

Additionally, if we are indeed living in McLuhan's global village, citizens of the techno-sphere, these Data-masks function as masks of the shaman. They are animistic deities brought out of the spirit world of the machine and into our material world, ready to tell us their secrets, or warn us of what's to come.

Bibliography

[1] Facebook, Last accessed December, 2014. http://www.facebook.com

[2] Alice Politics, "The Weekest Links: #stacktivism" *Stress.fm*, November 2013. Last accessed December 2014. http://blog.stress.fm/2013/11/the-weekest-links-stacktivism.html

[3] "FBI Next Generation Identification", Last accessed December 2014. http://www.fbi.gov/about-us/cjis/fingerprints_biometrics/ngi

[4] Jennifer Lynch, "FBI Plans to Have 52 Million Photos in its NGI Face Recognition Database by Next Year" *Eff.org*, October 2014. Last accessed December 2014. https://www.eff.org/deeplinks/2014/04/fbi-plans-have-52-million-photos-its-ngi-face-recognition-database-next-year

[5] Masahiro Mori, "The Uncanny Valley" *Ieee.org*, June 2012. Last checked December 2014. http://spectrum.ieee.org/automaton/robotics/humanoids/the-uncanny-valley

[6] Mark Fell, “Roadmap for The Emerging "Internet of Things"” *Carre Strauss*, 2014. Last accessed December 2014. http://carre-strauss.com/documents/IoT_Roadmap.pdf

[7] "FBI Biometric Center of Excellence", October 2014. Last accessed December 2014. http://www.fbi.gov/about-us/cjis/fingerprints_biometrics/biometric-center-of-excellence.

[8] Charles Darwin, *The Descent of Man, and Selection in Relation to Sex*. D. Appleton and Company, 1871.

[9] Francis Galton*, Inquiries into Human Faculty and its Development*. Macmillan, 1883.

[10] Timothy Lee, "Here's Everything We Know about Prism to Date " *Washington Post*, June 2013. Last accessed December 2014. http://www.washingtonpost.com/blogs/wonkblog /wp/2013/06/12/heres-everything-we-know-about-prism-to-date/

[11] James Bamford, "The NSA Is Building the Country's Biggest Spy Center (Watch What You Say)" *Wired*, March 2012. Last accessed December 2014. http://www.wired.com/2012/03/ff_nsadatacenter/all/1

[12] Kyle Chayka, "Biometric Surveillance Means Someone Is Always Watching", April 2014. Last accessed December 2014. http://www.newsweek.com/2014/ 04/25/biometric-surveillancemeans-someone-always-watching-248161.html

[13] Matt Stroud, "The Minority Report: Chicago's New Police Computer Predicts Crimes, but Is It Racist?" *The Verge*, February 2014. Last accessed December 2014. http://www.theverge.com/2014/2/19/5419854/the-minority-report-this-computer-predicts-crime-but-is-it-racist

[14] Facial Network, Last accessed December 2014. http://www.facialnetwork.com

[15] Quantified Self, Last accessed December 2014. http://www.quantifiedself.com

[16] Kashmir Hill, "Adventures in Self-Surveillance, aka The Quantified Self, aka Extreme Navel-Gazing" *Forbes*, October 2014. Last accessed December 2014. http://www.forbes.com/sites/kashmirhill/2011/04/07/adventures-in-self-surveillance-aka-the-quantified-self-aka-extreme-navel-gazing/

[17] Apple Watch, Last accessed December 2014. https://www.apple.com/watch/

[18] Yaniv Tiagman, Ming Yang, Marc' Aurelio Ranzato, Lior Wolf, "DeepFace: Closing the Gap to Human-Level Performance in Face Verification" *Facebook*. pages 1–8, April 2014.

[19] Kevin Kelly, "What Technology Wants" *Viking Press*, 2010.

[20] Brunor Latour, "Reassembling the Social: An Introduction to Actor-Network-Theory" *Oxford*, 2005.

[21] Raymond Kurzweil, "The Singularity Is Near: When Humans Transcend Biology" *Viking,* 2005.

[22] Ralph Abraham, Rupert Sheldrake, Terence K. McKenna, "The Evolutionary Mind: Trialogues at the Edge of the Unthinkable" *Trialouge Press*, 1998.

[23] Martin Heidegger, "Question Concerning Technology, and Other Essays" *Harper & Row*, 1982.

[24] Sterling Crispin, "CHARON: The Self and The Technological Other", 2013. http://www.sterlingcrispin.com/charon.pdf

[25] Marcos Novak, "Speciation, Transvergence,Allogenesis" pages 1–9, May 2002. www.mat.ucsb.edu/~marcos/transvergence.pdf

[26] Zach Blas, "Facial Weaponization Suite" 2014. Last accessed December 2014. http://www.zachblas.info/projects/facial-weaponization-suite/

[27] OpenCV, Last accessed December 2014. http://www.opencv.org/

[28] Matthew Turk, Alex Pentland, “Eigenfaces for Recognition” *Massachusetts Institute of Technology*, pages 1–16, June 2004.

[29] The AT&T Labratories Cambridge Database of Faces, 1994. Last accessed December 2014. https://www.cl.cam.ac.uk/research/dtg/attarchive/facedatabase.html.

[30] Adam Harvey, “CV Dazzle” 2010, Last accessed December 2014. http://ahprojects.com/projects/cv-dazzle/

[31] D.L. Williams, “Naval Camouflage 1914-1945: A Complete Visual Reference” *Naval Institute Press*, 2001.

[32] Christian Loclair Cedric Kiefer, Julia Laub, “Google Faces” 2013, Last Accessed December 2014. http://www.onformative.com/lab/googlefaces/

[33] Google Earth. Last accessed December 2014. http://www.google.com/earth/

[34] Greg Borenstein, “Machine Pareidolia: Hello Little Fella Meets Facetracker” January 2014, Last accessed December 2014. http://urbanhonking.com/ideasfordozens/2012/01/14/machinepareidolia-hello-little-fella-meets-facetracker/

[35] Roberto Brunelli, T. Poggio, “Face recognition: features versus templates” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol.15: pp.1042–1052. 1993.

[36] Rabia Jafri, Hamid R Arabnia, “A Survey of Face Recognition Techniques”

Journal of Information Processing Systems, 5(2):41–68, June 2009.

[37] Tazeem Choudhury, Brian Clarkson, Tony Jebara, Alex Petland, “Multimodal Person Recognition using Unconstrained Audio and Video” *MIT Media Laboratory*. 1998.

[38] Paul Viola, Michael Jones, “Rapid Object Detection using a Boosted Cascade of Simple Features*” Computer Vision and Pattern Recognition*, 2001.

[39] Yotam Abramson, Bruno Steux, “YEF Real-Time Object Detection” *Journal International Journal of Intelligent Systems Technologies and Applications.* Volume 2 Issue 2/3, June 2007 Pages 102-112.

[40] Processing. Last accessed December 2014. https://www.processing.org/

[41] The AlloSphere Research Facility: Software. Last accessed December 2014. http://mat.ucsb.edu/allosphere/software.php.

[42] OpenVDB, 2013. Last accessed December 2014. http://www.openvdb.org/.

[43] Sunflow The Open Source Render Engine. Last accessed December 2014. http://sunflow.sourceforge.net/.

[44] MeshLab. Last accessed December 2014. http://meshlab.sourceforge.net

[45] Bruce Lucas, Takeo Kanade, “An Iterative Image Registration Technique with an Application to Stereo Vision” *IJCAI'81 Proceedings of the 7th international joint conference on Artificial intelligence* - Volume 2 Pages 674-679. 1981.

[46] CCV A Modern Computer Vision Library. Last accessed December 2014.

http://libccv.org/.

[47] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller, "Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments" *University of Massachusetts, Amherst, Technical Report 07-49*, October, 2007. http://vis-www.cs.umass.edu/lfw/

[48] iPhoto, Last accessed December 2014. http://www.apple.com/mac/iphoto/

[49] Chang Huang, Haizhou Ai, Yuan Li, and Shihong Lao, "High-Performance Rotation Invariant Multiview Face Detection" *IEEE Transactions on Pattern Analysis and Machine Intelligence,* (Volume: 29, Issue 4) pages 671–686. 2007.

[50] Ben Jones, Lisa DeBruine, "Face Research" Last accessed December 2014. http://faceresearch.org/

[51] Roger Alsing, "Genetic Programming: Evolution of Mona Lisa", December 2008. Last accessed December 2014. http://rogeralsing.com/2008/12/07/genetic-programming-evolutionof-mona-lisa/

[52] Kevin Kelly, "The Three Breakthroughs That Have Finally Unleashed AI on the World" *Wired*, October 2014. Last accessed December 2014. http://www.wired.com/2014/10/future-of-artificial-intelligence

[53] Next Engine 3D Scanners. Last accessed December 2014 http://www.nextengine.com/

[54] Brett Lyons, "Additive Manufacturing in Aerospace: Examples and Research Outlook" *Spring Issue of The Bridge on Frontiers of Engineering* Volume 42, 2012.

[55] Deckard, C., "Method and apparatus for producing parts by selective sintering", U.S. Patent 4,863,538, filed October 17, 1986, published September 5, 1989.