Assessing Video Quality for Public Safety Applications Using Visual Acuity

Public Safety Communications Technical Report

May 2013

Homeland Security

Science and Technology

Initial Release: November 2012

Support to the Homeland Security Enterprise and First Responders: Office for Interoperability and Compatibility

Homeland Security
Science and Technology

Defining the Problem

Emergency responders—police officers, fire personnel, and emergency medical services personnel—need to share vital voice and data information across disciplines and jurisdictions to successfully respond to day-to-day incidents and large-scale emergencies. Unfortunately, for decades inadequate and unreliable communications have compromised their ability to perform mission-critical duties. Responders often have difficulty communicating when adjacent agencies operate on different radio bands, use incompatible proprietary systems and infrastructure, and lack adequate standard operating procedures and effective multi-jurisdictional, multi-disciplinary governance structures.

OIC Background

The U.S. Department of Homeland Security (DHS) established the Office for Interoperability and Compatibility (OIC) in 2004 to strengthen and integrate interoperability and compatibility efforts to improve local, tribal, state, and Federal emergency response and preparedness. Managed by the Science and Technology Directorate's Support to the Homeland Security Enterprise and First Responders Group, OIC helps coordinate interoperability efforts across DHS. OIC programs and initiatives address critical interoperability and compatibility issues. Priority areas include communications, equipment, and training.

OIC Programs

OIC programs address voice, data, and video interoperability. OIC is creating the capacity for increased levels of interoperability by developing tools, best practices, technologies, and methodologies that emergency response agencies can immediately put into effect. OIC is also improving incident response and recovery by developing tools, technologies, and messaging standards that help emergency responders manage incidents and exchange information in real time.

Practitioner-Driven Approach

OIC is committed to working in partnership with local, tribal, state, and Federal officials to serve critical emergency response needs. OIC's programs are unique in that they advocate a "bottom-up" approach. OIC's practitioner-driven governance structure gains from the valuable input of the emergency response community and from local, tribal, state, and Federal policy makers and leaders.

Long-Term Goals

Long-term goals for OIC include:

- Strengthen and integrate homeland security activities related to research and development, testing and evaluation, standards, technical assistance, training, and grant funding.
- Provide a single resource for information about and assistance with voice and data interoperability and compatibility issues.
- Reduce unnecessary duplication in emergency response programs and unneeded spending on interoperability issues.
- Identify and promote interoperability and compatibility best practices in the emergency response arena.

Intentionally Blank

Public Safety Communications Technical Report

Assessing Video Quality for Public Safety Applications Using Visual Acuity

DHS-TR-PSC-12-11
November 2012

Reported for: **The Office for Interoperability and Compatibility by the Public Safety Communications Research program**

Homeland Security

Science and Technology

Intentionally Blank

Publication Notice

Disclaimer

Contact Information

Please send comments or questions to: SandTFRG@hq.dhs.gov

Intentionally Blank

Contents

Figures

Tables

Abbreviations

3G	Third-Generation mobile telephony and data network (see CDMA and GSM)
CAVLC	Context-Adaptive Variable-Length Coding
CBR	Constant Bit Rate
CDMA	Code Division Multiple Access (see 3G and GSM)
CIF	Common Intermediate Format (352 x 288 pixels)
CODEC	Coding/Decoding
DHS	U.S. Department of Homeland Security
FirstNet	First Responder Network Authority
fps	frames per second
GOP	Group of Pictures
GSM	Global System for Mobile communications (see 3G and CDMA)
GUC	Generalized Use Class
kbps	kilobits per second
HD	High Definition
HRC	Hypothetical Reference Circuit
ITS	Institute for Telecommunication Sciences
ITU-T	International Telecommunication Union, Telecommunication Standardization Sector
LTE	Long-Term Evolution (see 4G)

LogMAR Logarithm of the Minimum Angle of Resolution

MHz Megahertz

MOS Mean Opinion Score

OIC Office for Interoperability and Compatibility

PPF pixels per foot

PSCR Public Safety Communications Research program

PSNR peak signal-to-noise ratio

S&T Science & Technology Directorate

TRV Target-Recognition Video

VGA Video Graphics Array (640 x 480 pixels)

VQEG Video Quality Experts Group

VQiPS Video Quality in Public Safety

Abstract

This report describes a laboratory study investigating the ability of human subjects to use a video system for various tasks. The particular tasks of interest focus on human targets and range from awareness of their presence to positive identification. The test simulates recorded video by allowing viewers significant control over how and when video sequences are displayed. The video sequences represent a variety of target sizes, motion, and lighting conditions. Various resolutions and encoder bit rates were used to allow us, public safety researchers at the Institute for Telecommunication Sciences (ITS), to recommend requirements for a particular task. Test subjects were asked to identify letters on an eye chart synthetically inserted into the video sequences. This allowed us to use visual acuity as a quality metric, which made it possible to separate measurements of the quality required for particular tasks from measurements of quality delivered by particular systems. The task-based subjective tests this report describes follow the test methods described in the International Telecommunication Union's Telecommunication Standardization Sector (ITU-T) Recommendation P.912 [1].

Key words: positive identification, video quality, subjective test methods, video, acuity

1 Introduction

Public safety practitioners use video for a wide variety of applications. These applications are often task-based and can be generalized as identifying targets (objects or persons of interest) in the video. While some applications rely on autonomous imaging systems to automatically perform these tasks, the majority of public safety video applications are intended for a human observer. Examples of public safety video applications include surveillance, telemedicine, urban search and rescue, and remote command and control.

For the applications listed above, subjective quality assessments such as mean opinion score (MOS) are not very useful. A useful measure of video quality for the public safety practitioner would be whether a viewer is able to recognize a target. If the operator of an unmanned urban search-and-rescue vehicle is not able to distinguish between an earthquake casualty and building debris on a video feed, then that video feed would be considered to be of poor quality and, therefore, less useful to that practitioner. The approach of this study is to evaluate video quality by measuring the success rates of test subjects who were assigned certain target-recognition tasks in a controlled environment.

Generalizing the factors that affect target-recognition success rates regardless of the application leads to results that may be useful to a wider variety of users in the public safety video system market. The success rate of target-recognition tasks depends on factors that are common to most public safety video applications; whether the recognition task is to be performed with live/real-time video or previously recorded video impacts the difficulty of

the target-recognition task. For this study, recognition tasks using a recorded video scenario were studied.

Categorizing public safety video applications in this way is the concept behind the generalized use class (GUC) framework, an integral part of the practitioner User Guide developed by the Video Quality in Public Safety (VQiPS) working group [2], [3]. The guide, *Defining Video Quality Requirements: A Guide for Public Safety*, provides recommendations for each GUC so that public safety video system consumers can match their specific video application to one of the GUCs and then use the provided recommendations to guide informed video system equipment purchases. Figure 1 illustrates the process by which an appropriate GUC for a specific public safety video application is identified. A user identifies whether his or her application involves large or small targets; high or low motion; bright, dim, or variable lighting; live or recorded video; and what discrimination level is required. By making these determinations, a user selects a specific GUC.

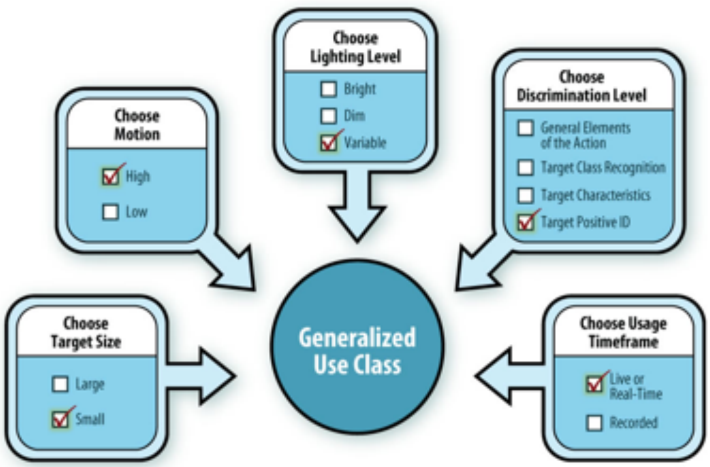

Figure 1. A representative example of determining an application's GUC

After identifying the appropriate GUC, the user is presented with a set of recommendations and resources linked to that GUC. Though the generated resources and recommendations may not exactly match the needs of each user's particular video application, the information provided will, at the very least, serve as a guide for informed purchasing decisions. This study was conducted in such a way that its results are consistent with both the User Guide [2] and GUC framework [3]. Further, the video system parameters used in this study contribute to the definition and understanding of target-recognition video (TRV) quality.

The U.S. Department of Homeland Security's (DHS) Office for Interoperability and Compatibility (OIC) has supported the Public Safety Communications Research (PSCR) program in providing content for the User Guide and GUC framework. More importantly, DHS OIC has enabled researchers at PSCR to perform laboratory experiments that explore the effects of network conditions on video quality for public safety applications.

This study explores how target-recognition tasks are affected by changing network conditions—specifically how target-recognition task success rates for H.264-compressed video sequences change when the resolution is reduced to Video Graphics Array (VGA; 640

x 480 pixels) or Common Intermediate Format (CIF; 352 x 288 pixels) levels and the encoder bit rate ranges from 64 kilobits per second (kbps) to 2048 kbps. The Coding/Decoding (CODEC) and resolutions used are representative of video that has been transmitted to mobile devices.

While video-based applications are commonly used on third-generation (3G) Code Division Multiple Access (CDMA) and Global System for Mobile communications (GSM) commercial networks, public safety users have been slow to embrace them. This is partly due to network constraints; most public safety wireless networks are not capable of supporting data or the bandwidths required by video-based applications. This is rapidly changing. On February 22, 2012, the First Responder Network Authority (FirstNet) was created when the Middle Class Tax Relief and Job Creation Act of 2012 [4] was enacted into law. FirstNet is charged with building, deploying, and operating a nationwide Long-Term Evolution (LTE)-based interoperable public safety broadband network.[1]

Commercial LTE has been in use since 2010 and differs from the proposed FirstNet LTE network in a number of fundamentally important ways. First, commercial LTE users enjoy 20 megahertz (MHz) of bandwidth per device, per channel, while public safety users are limited to 10 MHz of bandwidth per device, per channel. Additionally, commercial networks do not face the priority or scaling constraints that will be placed upon a nationwide public safety broadband network. In times of national emergency, public safety network traffic must be prioritized in such a way that the most important information gets to an incident commander, even when the network is overloaded. These constraints, combined with the already high bandwidth consumption of video applications, necessitate exploration of TRV quality under dynamic, LTE-like network conditions.

The test method used in this study adheres to ITU-T Recommendation P.912 [1], which addresses subjective assessment methods for video that is to be used for TRV tasks. A subjective test was performed using expert viewers and naïve viewers. The multiple-choice method was used (i.e., subjects were asked to recognize a target in a recorded video sequence and given several choices). Test-scene parameters were varied and five-bit rates were imposed on the H.264 encoder for each of two resolutions.

The tests reported here expand beyond ITU-T P.912 by introducing the idea of visual acuity. In this report, we attempt to use acuity as a single measure of video quality for public safety tasks. We are motivated by the fact that measuring video system requirements for each of the 96 GUCs separately would be prohibitively time consuming. We have used visual acuity to separate the dimension of discrimination level from the other characteristics that define a GUC. We performed one experiment (referred to as PS8) to measure the acuity delivered by a video system under a variety of lighting, motion, and target size conditions. We performed a separate experiment (referred to as PS9) to measure what level of acuity is required to achieve each discrimination level. When the acuity delivered by a video system equals or exceeds the acuity required for a particular discrimination level, we can recommend that system for a task associated with that discrimination level. While it would be more accurate to examine each GUC separately, this approach greatly reduced the time and effort required to make a complete set of recommendations.

[1] Source: http://www.ntia.doc.gov/category/public-safety

We believe acuity will be a particularly good measure of quality because the task of recognizing characters on a chart has many physical and psychological similarities to recognizing an object. For this reason we expect visual acuity to provide more relevant information than MOS, peak signal-to-noise ratio (PSNR), pixels per foot (PPF), and many other video quality metrics.

2 Targets and Scenario Groups

The test scenes used for this test also adhered to ITU-T Recommendation P.912, "Subjective Video Quality Assessment Methods for Recognition Tasks," [1]. This Recommendation introduces the concept of scenario groups. For the PS8 test, we used the same scenario groups as in the previous two experiments [5], [6]. Details are available in the reports, *Video Quality Tests for Object Recognition Applications* [5] and *Recorded-Video Quality Tests for Object Recognition Tasks* [6]. The scenario groups are designed to capture the three lighting levels, two motion levels, and two target sizes defined by the Video Quality in Public Safety (VQiPS) group. Although these parameters are not defined in a quantitative way, our scenes were designed to capture a variety of conditions. As in the previous tests, seven different objects served as targets: a gun, taser, police radio, cell phone, flashlight, mug, and soda can.

A new set of scenes was used for PS9. They were not designed to cover a wide range of GUCs. They were chosen because they presented human targets. A group of four male and four female actors were featured in this set of video sequences and at least one appeared in each sequence. This allowed us to ask subjects to perform a positive identification task within our multiple-choice framework. Viewers were also asked to perform other tasks involving human targets. Details can be found in Appendix A.

3 Acuity Charts

To measure visual acuity, we generated a set of reduced Logarithm of the Minimum Angle of Resolution (LogMAR) charts. These charts are roughly similar to the Snellen eye charts commonly used by physicians, but a few key differences make them recommended for any visual acuity research [7]. LogMAR charts are made up of rows of Sloan letters. A subset of the alphabet, C, D, H, K, N, O, R, S, V, and Z, is used. Each letter is randomly selected with a uniform distribution. Each row is made up of characters of a particular size. At the bottom of the chart, the smallest row is sized so that the height of each character will be exactly five pixels when the chart has been resized to fit within a VGA frame. This is regarded as the minimum number of pixels necessary to reliably recognize characters with good contrast and perfect quality. For this reason, there is no purpose in using any smaller characters. Each row uses characters 1.414 times the height of the row directly beneath it so that the size of the characters doubles every two rows. Eight rows were used for each chart. Typically, a LogMAR chart includes five characters on each line, but our tests used only three characters per line. This is because we are not interested in measuring the visual acuity of particular viewers. We can combine results from multiple viewers to measure the acuity of the group and save time for individual viewers by reducing the number of letters in each row. Figure 2 shows examples of the charts used in PS8 and PS9.

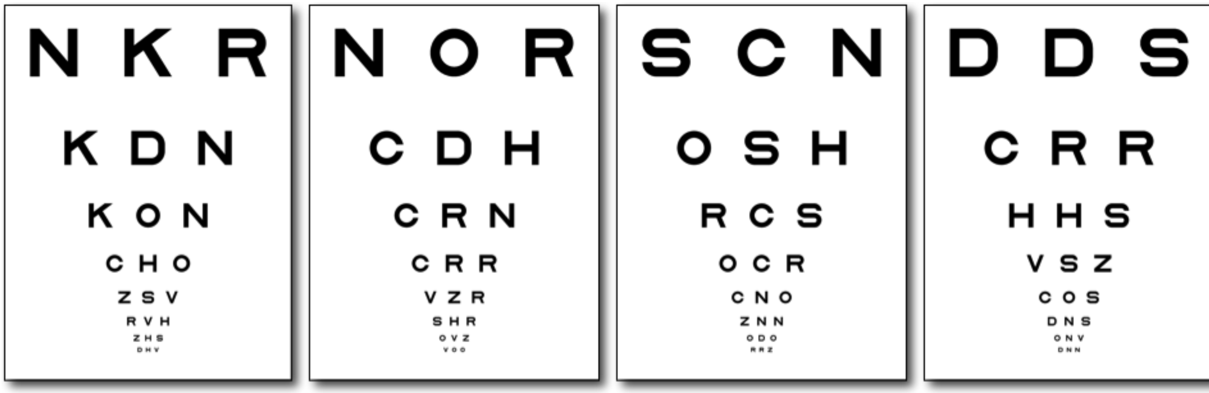

Figure 2. An example of acuity charts

4 Processed Scenes

All high definition (HD) clips were down-converted to two display resolutions: VGA and CIF. The frame rate was kept constant at 29.97 frames per second (fps).

Each clip then had an acuity chart synthetically inserted into it. For PS8, the chart moved across the screen at the same rate (in pixels per frame) as the target object. The contrast of the chart was also adjusted to reflect the lighting on the object. This changed over time for variable lighting conditions. For the moving charts, the shutter speed of the camera was simulated by averaging the light that the lens would capture over an entire 1/30th of a second. This matches our observations of the workings of practical video cameras. For PS9, we were not attempting to control motion or lighting, so each of the charts was stationary and nothing was done to affect the contrast of the chart. The chart images were generated at very high resolution. For each test, great care was taken to ensure that the charts were resized accurately, taking any simulated motion into account.

The clips were then compressed via H.264 encoding at various bit rates. Five bit rates were chosen for each resolution. The bit rates were chosen to represent a wide range of resultant video quality and to represent a wide range of bandwidth requirements. Table 1 lists bit rates.

Table 1. Encoder Bit Rates

Resolution	**Bit Rates (kbps)**
CIF	64, 128, 256, 512, 1024
VGA	128, 256, 512, 1024, 1536

Each combination of resolution and bit rate is what is referred to as a Hypothetical Reference Circuit, or HRC. The Video Quality Experts Group, or VQEG, uses this term. HRC refers to the distortion to the video signal that is being tested—in this particular case, combinations of compression and resolution reduction.

After being processed through an HRC, each clip was decoded and resampled to VGA resolution. By doing this, we minimized the computational requirements for our test computer to display the sequences, although this required faster storage. Appendix B provides further details.

5 PS8 Test Design

5.1 Test size

As in our previous two experiments ([5] and [6]), there were 14 scenario groups. Seven objects were used in twelve of the groups, and only six were used in the other two. This results in a total of 96 source scenes. A different random chart was inserted into each source to avoid any memorization effects. After processing through 10 different HRCs, there were 960 total clips. To avoid showing any viewer the same source twice, each viewer only saw one HRC for a given source. The HRCs were chosen so that the distribution of HRCs among viewers was as uniform as possible. This meant that each viewer saw 96 different video sequences over the course of the test. This generally resulted in a total test time of 60 to 90 minutes. Additionally, at the beginning of the test, each viewer saw a training sequence showing each target object with a label; they then took a practice test consisting of four additional clips for familiarization with the test software.

5.2 Viewers

Thirty-nine viewers participated in the test. In accordance with ITU-T Recommendation P.912 [1], expert viewers were recruited. These viewers had experience as practitioners in law enforcement, the fire service, or emergency medical services.

Viewers were screened for visual acuity and color vision by way of Snellen and Ishihara tests, respectively. Viewers were not automatically excluded from the test if they demonstrated impaired acuity or color vision. In our previous tests ([5] and [6]), an analysis revealed performance of the recognition task was not significantly affected by such visual impairments.

5.3 Test Environment and Software

Viewing conditions generally followed the recommendations in ITU-T P.910 [8]. One exception was that the viewers could choose their viewing distance, and it was not recorded. Viewing distance is measured relative to the height of the picture being displayed. It is reasonable to assume that viewers' chosen viewing distances most likely fell into the recommended range of one to eight times the picture height, given an approximate picture height of five inches.

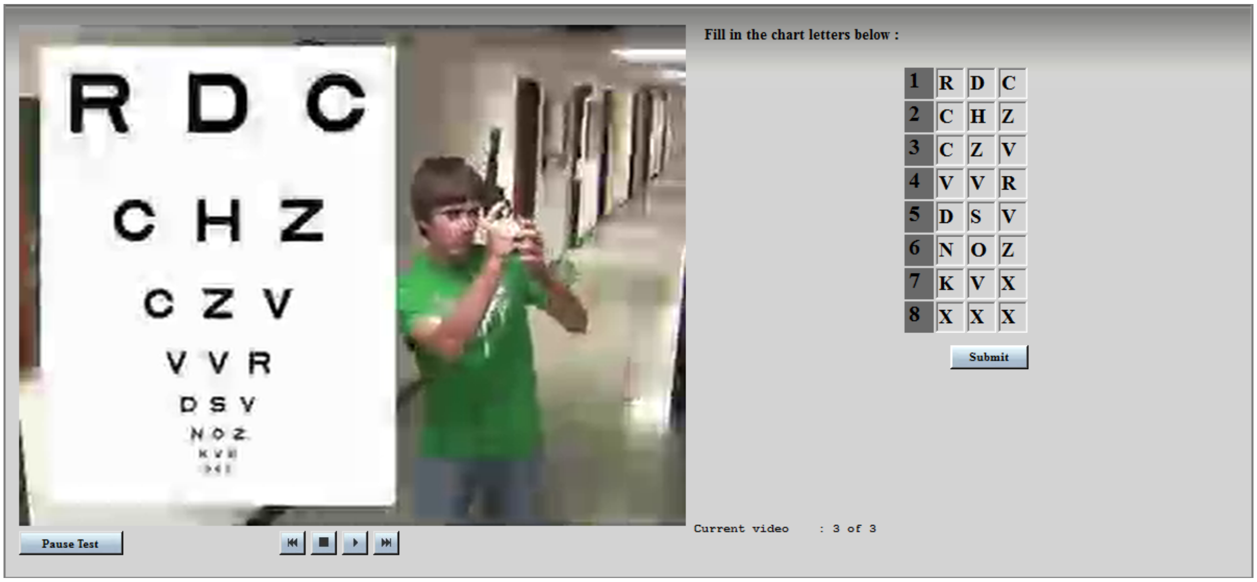

Figure 3. Test software user interface with acuity chart input

Figures 3 and 4 illustrate the user interface for the test. Viewers were shown processed video clips, asked to identify letters on the acuity chart, and asked specific questions about the target of interest. Viewers were allowed to view the clip as many times as they chose; they also had the option to rewind and fast-forward, pausing and advancing the video frame by frame as much as they wished. This degree of control allowed viewers to find the individual frame that was most useful to the task at hand. The user interface recorded all such interactions with the software for future analysis.

Figure 4. Test software user interface with multiple-choice input

The test length was expected to be approximately 90 minutes spent actively viewing videos and providing answers to the target-recognition questions. Because the viewers controlled

their own interaction with the video, the total time for the test could vary widely among viewers. Subjects were free to take breaks as needed. Appendix D shows the viewer instructions.

6 PS9 Test Design

6.1 Test size

There were 48 source scenes and 10 HRCs under test. Eight different actors each appeared in six scenes each. The total number of clips when all source scenes were processed with all HRCs was 480. To reduce test length and viewer fatigue, each viewer did not see each clip, but instead saw only one HRC for each source. Additionally, at the beginning of the test, each viewer took a practice test consisting of three additional clips for familiarization with the test software. These practice clips included charts that did not appear later in the test, so there was no risk of memorization.

6.2 Viewers

Forty-five viewers participated in the test. In accordance with ITU-T Recommendation P.912 [1], expert viewers were recruited. These viewers had experience as practitioners in law enforcement, the fire service, or emergency medical services. Naïve viewers were also used. These viewers were recruited from a temporary employment agency and were paid for their time. Based on previous experiments, we believe these two different sets of viewers produce equivalent results. For the purposes of this report, we do not distinguish between the two sets of viewers, but this remains an important subject for future analysis.

Viewers were screened for visual acuity and color vision by way of Snellen and Ishihara tests, respectively. Viewers were not automatically excluded from the test if they demonstrated impaired acuity or color vision. In our previous tests ([5] and [6]), an analysis revealed performance of the recognition task was not significantly affected by such visual impairments.

6.3 Test Environment and Software

Viewing conditions generally followed the recommendations in ITU-T P.910 [8]. One exception was that the viewers could choose their viewing distance, and it was not recorded. However, it is reasonable to assume that viewers' chosen viewing distances most likely fell into the one to eight picture heights recommendation, given an approximate picture height of five inches.

Figures 3 and 4 illustrate the user interface for the test. As in PS8, viewers were shown processed video clips, asked to identify letters on the acuity chart, and asked to answer specific questions about the target of interest. For each scene, there were four multiple-choice questions. Viewers were first asked a question about how many people were present in the scene. This was intended to measure the viewers' ability to perform a task at the "general elements" discrimination level. Second, viewers were asked whether the target actor in the scene was male or female. This constituted a task at the "classification"

discrimination level. Third, viewers were asked to indicate whether the target actor was wearing a hat, glasses, a necklace, some combination of the three, or none of the above. This was designed as a task representing the "characteristics" discrimination level. Finally, each viewer was asked to perform a "positive identification" task by selecting each target actor from a choice of eight people. Labeled pictures of the eight actors' faces were displayed for the viewers throughout the test. See Appendix A for source scenes.

Viewers were allowed to view the clip as many times as they chose; they also had the option to rewind and fast-forward, pausing and advancing the video one frame at a time as much as they wished. The user interface recorded all such interactions with the software for future analysis.

The test length was expected to be approximately 90 minutes spent actively viewing videos and providing answers to the multiple-choice questions. While there were fewer clips than PS8, more questions were asked about each clip, so the total viewing time was similar. Because the viewers controlled their own interaction with the video, the total time for the test could vary widely among viewers. Subjects were free to take breaks as needed. Appendix D lists viewer instructions.

7 Results

For PS8, each viewer's responses to each clip were examined and compared against the ground truth for that clip. The total number of characters correctly recognized in each row of the acuity chart (ranging from zero to three) was tallied. We also determined whether each viewer correctly recognized the object they were shown for each clip. We then treated all viewers as statistically equivalent and combined their results for each clip. We assumed that the object recognition task is essentially the same, regardless of which of the seven objects is shown, and combined all the results across objects. For each combination of scenario group and HRC, we were then left with a total number of times viewed, a total number of correct object recognitions, and a total number of characters recognized for each of the eight sizes. Appendix E details these results. By dividing the number of correct object recognitions by the total times each combination of scenario group and HRC were shown, we calculated the object recognition rates.

The detailed results of PS9 are too voluminous to be included in this report, but they are available upon request. We combined results from different viewers assuming they are statistically equivalent. For each video clip (each combination of source and HRC), we totaled the number of times the clip was shown, the number of times acuity chart characters of each size were recognized, the number of times the correct number of people were identified in the clip, and the number of times the gender of the target actor was correctly identified. We also totaled the number of times that viewers correctly indicated whether the viewer was wearing a hat, wearing glasses, or wearing a necklace. We separately totaled the number of times a viewer answered all three elements of this question correctly. By analyzing the three aspects of this question separately, we have laid the foundation for a future analysis of what aspects of this question were most difficult and how close the viewers were to answering correctly when their answers were not perfect. Finally, we totaled the number of times viewers correctly identified the target actor for each

clip. By dividing each of these totals by the total time each clip was shown, we were able to calculate the rate at which each task was successfully performed.

Because PS9 was not designed with meaningful scenario groups (as PS8 was), we were able to combine the results from every source clip and examine how each of the 10 HRCs produced changes in visual acuity as well as the ability to perform particular tasks. Appendix F provides these totals.

7.1 Measuring Acuity

The exact definition of visual acuity we have chosen to use is the inverse of the height (measured in pixels) of the smallest reliably-recognizable characters on the acuity charts. For the purposes of this report, we chose a 90-percent recognition rate as the minimum to be considered reliable. In these experiments, this means that if 10 characters with a height of 5 pixels are shown, and 9 or 10 of them are correctly recognized, the acuity is 0.2. If fewer than 9 are recognized, the acuity must be lower. In addition, if fewer than 9 characters with a height of 7.1 pixels are correctly recognized, but 9 or 10 characters with a height of 10 pixels are recognized, the acuity is 0.1.

In this way, we have developed a metric based upon the acuity charts used in our video clips that increases as recognition ability increases. This metric also appropriately deals with the statistical nature of these recognition tasks.

With this approach, we were able to calculate acuity values based on the totals of correctly-recognized characters for each scenario group and HRC combination in PS8 and for each clip in PS9.

8 Conclusions

The purpose of the PS9 test was to calculate acuity requirements for particular tasks. With these requirements we can use the PS8 data to recommend appropriate video systems. To come up with these acuity requirements, we first define the tasks in mathematical terms and then use the data in Appendix F to find an acuity level that can be considered sufficient.

We defined the positive identification task as requiring viewers to correctly identify the target actor from among 8 choices 90 percent of the time. Our PS9 results show that this is achieved in systems delivering an acuity of 0.144.

For the target characteristics task, we required the viewers to perfectly describe whether the target actor was wearing a hat, glasses, or a necklace 75 percent of the time. It should be noted that because an error in any of the three elements of this task will fall short of a perfect description, it is not unreasonable to have a somewhat lower reliability. Our PS9 results show that this task was achieved at an acuity of 0.1.

For the classification task, we asked viewers to identify the gender of the target actor. This task was performed with 98- to 99-percent accuracy with every HRC. Essentially, this means that the task was too easy or that the test was not designed with sufficiently degraded video to cause viewers to fail at this task. Essentially, the minimum quality

required for the classification task is outside of the range of what was measured in this test. The lowest acuity measured in the test was 0.07. What we can say is that an acuity of 0.07 should be sufficient for the recognition task. We will proceed to make recommendations on that basis.

For the "general elements" task, the rate of correct answers hovered around 90 percent regardless of HRC, acuity level, or quality. In fact, viewers seemed to perform the task slightly better when acuity was lower. This indicates that video quality was not the limiting factor for the performance of this task. Searching the entire video sequence for any time a person appears and successfully counting every person in the scene is more cognitively demanding than the other tasks in this test. Additional testing would be required to define a "general elements" task that is more dependent on video quality; however, such a task is likely to be easy if video is being transmitted at all and only difficult when quality is clearly unacceptable. For this reason, we believe it is reasonable to make recommendations based on an acuity level slightly lower than anything actually measured in this test. We have chosen to set an acuity requirement of 0.05 for the "general elements" task.

Table 2. Acuity Requirements

Discrimination Level	Acuity Required
General Elements	0.05
Classification	0.07
Characteristics	0.1
Positive Identification	0.144

We have now established acuity requirements for each discrimination level as Table 2 shows. Using data from PS8, we can determine which HRCs meet these requirements under a variety of lighting and motion conditions. Because target size is not defined in any quantitative way, we have chosen to always recommend our higher resolution (VGA) for small targets and our lower resolution (CIF) for large targets. Previous testing [6] has indicated that live and recorded video do not produce significantly different results for recognition tasks. Using these assumptions, all that remains is to find the lowest bit rates that provide acceptable acuity for a particular GUC. Table 3 shows these results.

8.1 Recommendations

Table 3 shows the bit rate and resolution recommendations we are making for each of the VQiPS GUCs. Some of the bit rate recommendations are marked with an asterisk. This indicates that no bit rate tested produced sufficient acuity to perform the desired tasks. In these cases, we have recommended the maximum bit rate tested. Based on earlier results [5], we believe these are situations in which no amount of additional bandwidth will produce the desired quality. These are generally cases involving very poor lighting and possibly high motion. In these circumstances, the compression of the video is not limiting the quality and the only way to see an improvement is to improve the lighting.

Nevertheless, if poor lighting or high motion is unavoidable, we recommend high bit rates to ensure that the CODEC will not further impair video quality.

Table 3. Summary of Recommendations

Scenario	General Elements	Classification	Characteristics (75-percent perfect)	Positive ID (90-percent accuracy)
Bright light, low motion, large target	64 kbps, CIF	64 kbps, CIF	64 kbps, CIF	128 kbps, CIF
Bright light, low motion, small target	128 kbps, VGA	128 kbps, VGA	128 kbps, VGA	128 kbps, VGA
Bright light, high motion, large target	64 kbps, CIF	64 kbps, CIF	128 kbps, CIF	*1024 kbps, CIF
Bright light, high motion, small target	128 kbps, VGA	128 kbps, VGA	128 kbps, VGA	256 kbps, VGA
Dim light, low motion, large target	128 kbps, CIF	128 kbps, CIF	256 kbps, CIF	512 kbps, CIF
Dim light, low motion, small target	256 kbps, VGA	256 kbps, VGA	512 kbps, VGA	512 kbps, VGA
Dim light, high motion, large target	128 kbps, CIF	256 kbps, CIF	512 kbps, CIF	*1024 kbps, CIF
Dim light, high motion, small target	512 kbps, VGA	512 kbps, VGA	1024 kbps, VGA	*2048 kbps, VGA
Variable light, low motion, large target	256 kbps, CIF	512 kbps, CIF	512 kbps, CIF	*1024 kbps, CIF
Variable light, low motion, small target	256 kbps, VGA	1024 kbps, VGA	1024 kbps, VGA	*2048 kbps, VGA

Scenario	General Elements	Classification	Characteristics (75-percent perfect)	Positive ID (90-percent accuracy)
Variable light, high motion, large target	256 kbps, CIF	1025 kbps, CIF	*1024 kbps, CIF	*1024 kbps, CIF
Variable light, high motion, small target	512 kbps, VGA	512 kbps, VGA	*2048 kbps, VGA	*2048 kbps, VGA

8.2 Limitations and Future Work

A significant limitation of the PS8 and PS9 experiments is the tasks that have been studied. While we believe these tasks are representative of many public safety applications, we have no basis for generalizing beyond the specific tasks we asked our viewers to perform. However, by introducing the idea of visual acuity as a quality metric, we have created a framework to run experiments for any task that may become relevant in the future.

Many opportunities for further analysis of the data presented here have already been suggested. In addition, significant work could be done with this data to evaluate the utility of acuity as a quality metric. Figure 5 shows a scatter plot of the object recognition rate versus visual acuity for the PS8 data. There appears to be some relationship between the two, but acuity only captures a part of what is happening to the utility of a video system as the quality changes. Further work should be done to determine when acuity is a meaningful metric and when it is not. Different quality metrics could also be explored to attempt to find a better option.

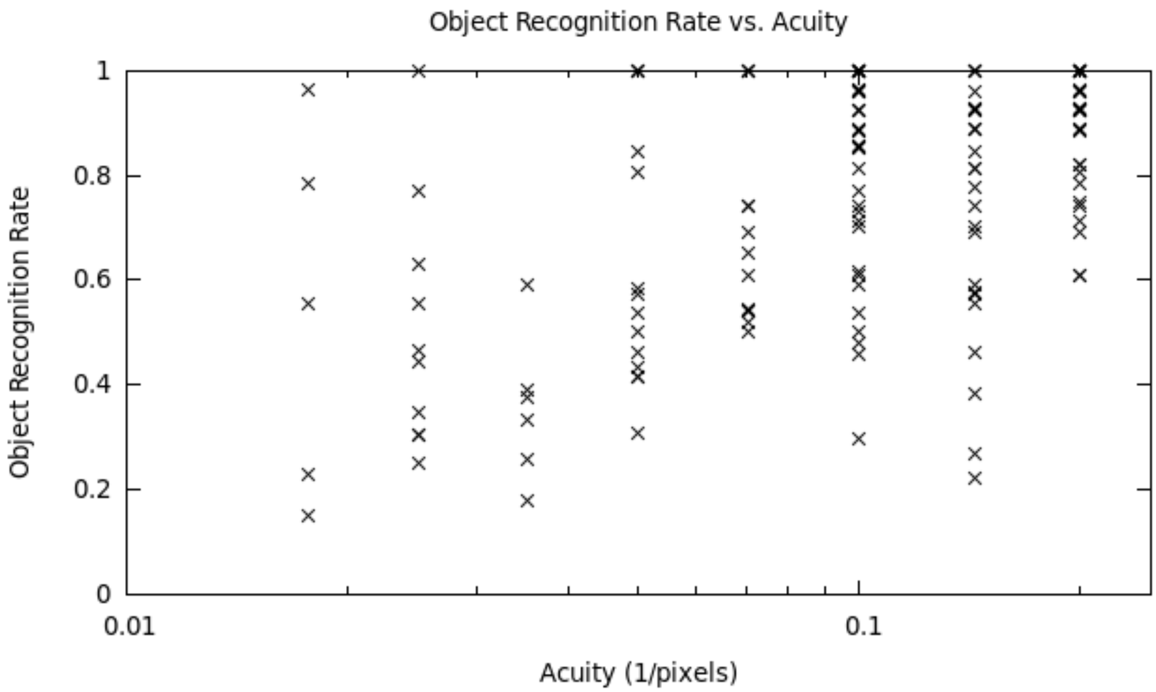

Figure 5. PS8 Data: Object Recognition Rate vs. Acuity

Additionally, there are significant opportunities for future work related to the new public safety broadband network. The tests we have performed for this report and previous ones [5] [6] have focused on measuring the relationship between resources devoted to digital video (bit rate or bandwidth) and the utility of that video. This kind of information readily lends itself to optimization algorithms. Using the results of these studies, we could potentially calculate the optimal distribution of resources over a network as demands increase from multiple users and multiple video streams. This is particularly practical for public safety in emergency situations.

9 References

[1] ITU-T Recommendation P.912, “Subjective Video Quality Assessment Methods for Recognition Tasks,” Geneva, 2008 (available: http://www.itu.int), Cited September 2012.

[2] *Defining Video Quality Requirements: A Guide for Public Safety*, Volume 1.0, July 2010 (available: http://www.safecomprogram.gov/SiteCollectionDocuments/3aVideoUserRequirementGuidedoc.pdf), Cited September 2012.

[3] “Guide to Defining Video Quality Requirements” (available: http://www.pscr.gov/outreach/vqips/vqips_guide/define_vid_qual_reqs.php), Cited September 2012.

[4] United States Congress, H.R. 3630 “Middle Class Tax Relief and Job Creation Act of 2012,” Title VI (the Spectrum Act) (P.L. No. 112-96 §6601), February 22, 2012.

[5] "Video Quality Tests for Object Recognition Applications," September 2010 (available: http://www.safecomprogram.gov/library/Lists/Library/Attachments/231/Video_Quality_Tests_for_Object_Recognition_Applications.pdf), Cited 2012.

[6] "Recorded-Video Quality Tests for Object Recognition Tasks," September 2011 (available: http://www.pscr.gov/outreach/safecom/vqips_reports/RecVidObjRecogn.pdf), Cited September 2012.

[7] Grosvenor, Theodore (2007). *Primary Care Optometry*. St. Louis, Missouri: ELSEVIER. pp. 174–175.

[8] ITU-T Recommendation P.910, "Subjective Video Quality Assessment Methods for Multimedia Applications," Geneva, 1999 (available http://www.itu.int), Cited September 2012.

Intentionally Blank

Appendix A Source Scenes

The test scenes used for this test followed ITU-T Recommendation P.912 [1]. For PS9, six scenes were chosen for each of eight actors. Four male actors and four female actresses were used. Throughout the test, the image in Figure 6 was on display as a reference to the viewers. The image displays a photo of each actor along with an identifier (name).

Figure 6. Test targets, as shown in background during each test

A representative sample of still frames from PS9 is shown in the next several figures. Scenarios included simulated bank robberies, simulated shoplifting, and general surveillance.

Figure 7. Frame from an original clip

Figure 8 shows the same frame as Figure 7 but is cropped to 4:3 and includes an acuity chart (see Figure 2).

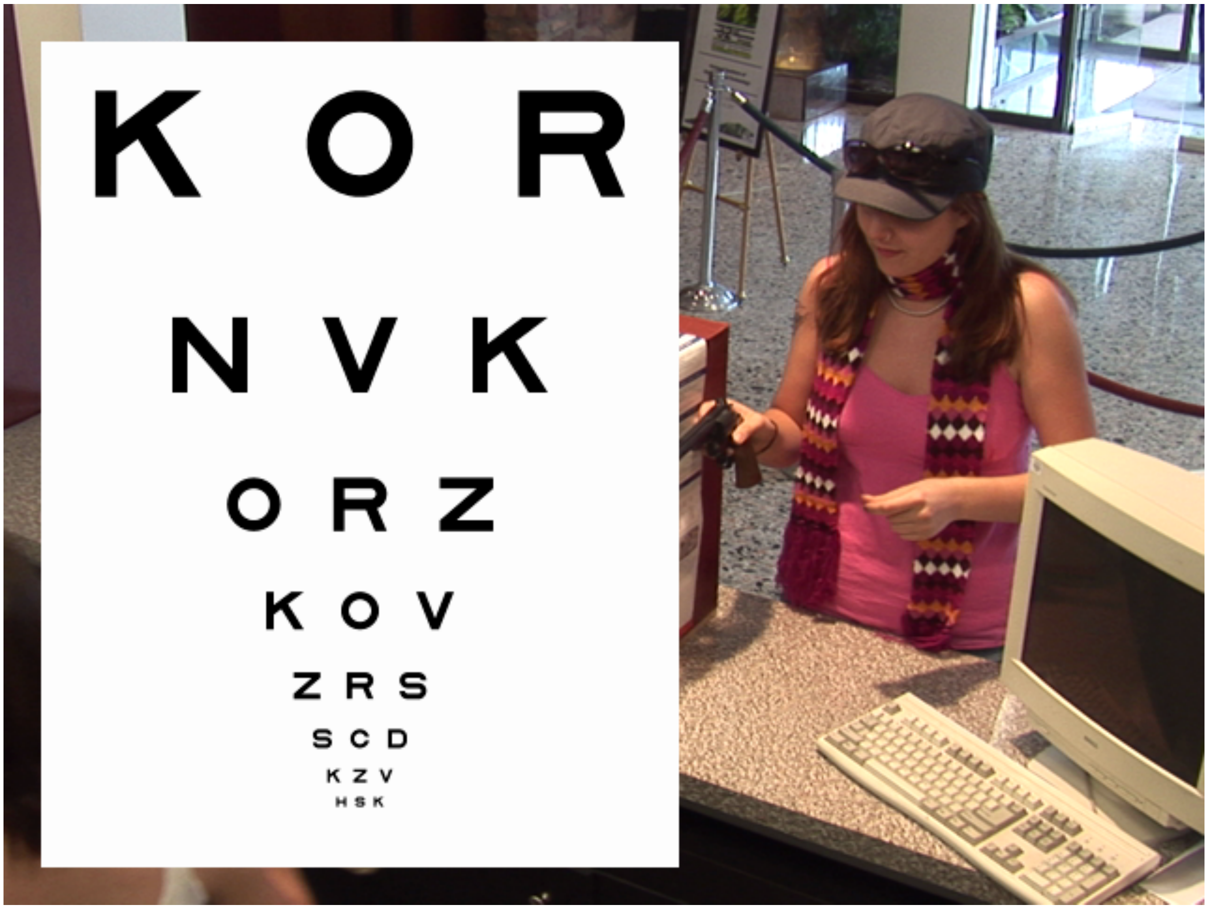

Figure 8. Frame cropped to 4:3 with acuity chart

Figure 9. Target walking hallway, unprocessed video

Figure 10 shows the same frame as Figure 9 but is down-converted to CIF at 512 kbps.

Figure 10. Target walking hallway, video down-converted to CIF at 512 kbps

Figure 11 shows the same frame as Figure 10, but is down-converted to CIF at 64 kbps.

Figure 11. Target walking hallway, video down-converted to CIF at 64 kbps

Intentionally Blank

Appendix B Processed Scenes

All HD clips were down-converted to two display resolutions: VGA (640 x 480 pixels) and CIF (352 x 288 pixels). The CIF resolution clips were enlarged to cover the same area on the screen as the VGA clips. The interpolation was done with a Lanczos filter. The frame rate was kept constant at 29.97 fps.

The clips were then compressed via H.264 encoding at various bit rates. Five bit rates were chosen for each resolution. The bit rates were chosen to represent a wide range of resultant video quality. Table 4 lists bit rates.

Table 4. Encoder Bit Rates

Resolution	Bit Rates (kbps)
CIF	64, 128, 256, 512, 1024
VGA	128, 256, 512, 1024, 1536

Encoding was done with TMPGEnc Xpress 4.0 software, which employed the MainConcept H.264 encoder. Table 5 lists the software settings used.

Table 5. Software Settings for H.264 Encoding

Parameter	Setting
Profile	Baseline
Level	Automatic
Frame Rate	29.97 fps
Bit rate mode	One-pass CBR (Constant Bit Rate)
Motion Search Range	63
Detect Scene Changes?	Yes
GOP (Group of Pictures) Length	33
B-Frame Count	0
Quantization Parameters	I Picture: 24 P Picture: 25
Entropy Coding Mode	CAVLC (Context-Adaptive Variable-Length Coding)

Parameter	Setting
Motion Estimation Sub-pixel Mode	Quarter-pixel

Intentionally Blank

Appendix C Notes on Experimental Design

Randomization

The total number of clips in PS8 when all source scenes were processed with all HRCs was 960. The total number of clips in PS9 was 480. To reduce test length, each viewer did not see each clip, but instead saw only one HRC for each source scene. This reduced viewer fatigue and any memorization effects. This resulted in each PS8 viewer seeing 96 clips and each PS9 viewer seeing 48 clips.

The clips to be viewed were selected in advance and distributed uniformly among the viewers. The order in which the clips were presented was randomized for each viewer in advance of the test.

Intentionally Blank

Appendix D Viewer Instructions

PS8 Instructions

Public Safety Video Quality Test
Overview for Subjects

Thank you for participating in our study. This study concerns the quality of video images for use in public safety applications. As a likely user of next-generation devices for public safety applications, we are interested in whether the videos to be presented are of sufficient quality to be used by you to perform several different potential tasks.

The study examines video used in a **recorded** situation and the ability to use this video to make decisions on how to respond to an incident. This study applies to video that has been recorded for later examination. The task currently under investigation is gathering information about people appearing in video scenes, from awareness of their presence to positive identification. You will be asked to answer specific questions regarding content in the video.

Each scene is approximately 10 seconds long. While the clip is playing, you may pause or step backward or forward frame by frame. You may replay each clip as many times as you wish. You will then be asked to type in the letters you saw on an eye chart in the video and answer a series of multiple-choice questions about the video. The test software will record your answers, as well as when you paused, replayed, or stepped through frames of the clip, and the total time you spent on each clip.

Important Clarifications for Multiple-Choice Questions
Some multiple-choice questions ask how many people appear in a certain part of a scene. If any part of a person appears in the given part of the scene, they should be counted, even if you can't see their face or if they later leave the scene.

Other questions ask if a person is wearing particular items of clothing or jewelry. Be aware that people can wear these items in unconventional ways. For example, if a person is wearing sunglasses on top of a hat, you should consider that person to be wearing sunglasses even though the sunglasses are not covering their eyes.

Text Entry and Multiple Choice Instructions
If you are unable to recognize a letter on the chart, enter an "X" in the corresponding box. Only enter "X" if you are completely unsure. If you recognize any characteristics of the letter, please make your best guess. Then, please choose the object that best matches what you saw in the video. For this question, there is no "other" or "I don't know" option. Therefore, please select the answer you believe to be most likely.

You will be asked to participate in one viewing session that is approximately 90 minutes long. A practice session will be presented to help you become familiar with the scene

material and rating process, as well as a clip showing the objects you might see in the videos. You may take a break at any time during the session.

PS9 Instructions

Public Safety Video Quality Test
Overview for Subjects

Thank you for participating in our study. This study concerns the quality of video images for use in public safety applications. As a likely user of next-generation devices for public safety applications, we are interested in whether the videos to be presented are of sufficient quality to be used by you to perform several different potential tasks.

The study examines video used in a **recorded** situation and the ability to use this video to make decisions on how to respond to an incident. This study applies to video that has been recorded for later examination. The task currently under investigation is gathering information about people appearing in video scenes, from awareness of their presence to positive identification. You will be asked to answer specific questions regarding content in the video.

Each scene is approximately 10 seconds long. While the clip is playing, you may pause or step backward or forward frame by frame. You may replay each clip as many times as you wish. You will then be asked to type in the letters you saw on an eye chart in the video and answer a series of multiple-choice questions about the video. The test software will record your answers, as well as when you paused, replayed, or stepped through frames of the clip, and the total time you spent on each clip.

Important Clarifications for Multiple-Choice Questions
Some multiple-choice questions ask how many people appear in a certain part of a scene. If any part of a person appears in the given part of the scene, they should be counted, even if you can't see their face or if they later leave the scene.

Other questions ask if a person is wearing particular items of clothing or jewelry. Be aware that people can wear these items in unconventional ways. For example, if a person is wearing sunglasses on top of a hat, you should consider that person to be wearing sunglasses even though the sunglasses are not covering their eyes.

Text Entry and Multiple Choice Instructions
If you are unable to recognize a letter on the chart, enter an "X" in the corresponding box. Only enter "X" if you are completely unsure. If you recognize any characteristics of the letter, please make your best guess. Then, please choose the object that best matches what you saw in the video. For this question, there is no "other" or "I don't know" option. Therefore, please select the answer you believe to be most likely.

You will be asked to participate in one viewing session that is approximately 90 minutes long. A practice session will be presented to help you become familiar with the scene

material and rating process, as well as a clip showing the objects you might see in the videos. You may take a break at any time during the session.

Intentionally Blank

Appendix E Detailed PS8 Results

Table 6. Detailed PS8 Result Data

SG	HRC	Times shown	Row1	Row2	Row3	Row4	Row5	Row6	Row7	Row8	Obj Rec
IBL	cif0064	28	81	70	51	14	1	0	0	0	27
IBL	cif0128	26	78	77	77	77	70	39	3	1	26
IBL	cif0256	27	81	81	81	81	80	69	25	0	27
IBL	cif0512	26	78	78	78	78	78	73	27	3	26
IBL	cif1024	26	78	78	78	78	77	75	34	3	26
IBL	vga0128	28	83	82	75	30	1	0	0	0	28
IBL	vga0256	27	81	81	81	79	74	40	9	0	27
IBL	vga0512	28	83	81	81	79	81	78	52	7	27
IBL	vga1024	27	81	81	81	80	81	79	52	14	27
IBL	vga2048	26	78	78	78	77	77	76	58	15	25
IBR	cif0064	28	80	69	45	9	2	0	0	0	22
IBR	cif0128	28	84	83	83	83	55	22	0	0	28
IBR	cif0256	26	78	77	76	77	76	56	10	0	26
IBR	cif0512	27	81	81	79	81	80	76	21	1	27
IBR	cif1024	28	84	84	84	83	82	76	28	0	28
IBR	vga0128	26	78	75	66	21	5	0	0	0	20
IBR	vga0256	28	84	84	82	78	58	38	6	1	28
IBR	vga0512	26	78	78	78	77	78	76	21	4	25
IBR	vga1024	26	78	78	75	78	77	72	26	1	26
IBR	vga2048	28	84	84	84	84	84	79	33	6	28
IDL	cif0064	23	68	63	34	2	2	1	0	0	7
IDL	cif0128	21	63	63	62	45	16	1	0	0	7
IDL	cif0256	24	71	72	72	72	62	17	0	0	14
IDL	cif0512	24	71	72	71	72	65	44	3	0	13
IDL	cif1024	20	59	60	60	60	57	46	12	1	13
IDL	vga0128	24	71	71	66	36	2	0	0	0	9
IDL	vga0256	23	67	69	66	67	26	5	0	1	10
IDL	vga0512	24	72	71	72	72	65	36	2	0	12
IDL	vga1024	22	66	65	66	66	66	62	25	3	13
IDL	vga2048	23	69	67	64	66	66	59	20	0	14
IDR	cif0064	22	64	57	30	2	0	0	0	0	5

SG	HRC	Times shown	Row1	Row2	Row3	Row4	Row5	Row6	Row7	Row8	Obj Rec
IDR	cif0128	23	69	69	64	40	7	0	0	0	9
IDR	cif0256	24	71	71	71	69	50	25	1	0	10
IDR	cif0512	22	66	65	63	63	54	37	2	0	11
IDR	cif1024	24	71	72	70	72	71	58	11	0	13
IDR	vga0128	23	68	64	56	44	6	1	0	0	8
IDR	vga0256	23	69	64	61	48	7	0	1	2	7
IDR	vga0512	22	65	66	63	66	61	50	2	0	12
IDR	vga1024	24	71	72	71	68	64	58	23	0	10
IDR	vga2048	24	72	69	72	72	72	65	36	7	11
IDS	cif0064	24	63	69	47	12	2	0	0	0	6
IDS	cif0128	28	84	81	80	54	33	3	0	0	5
IDS	cif0256	26	77	78	78	74	57	34	1	0	12
IDS	cif0512	27	81	81	81	79	78	73	37	4	13
IDS	cif1024	26	78	78	77	78	72	75	47	13	16
IDS	vga0128	27	81	79	77	66	11	0	0	0	7
IDS	vga0256	26	78	77	77	72	41	22	9	0	8
IDS	vga0512	28	84	84	84	84	74	73	40	7	16
IDS	vga1024	28	84	84	84	84	83	78	67	42	14
IDS	vga2048	26	77	78	78	78	77	77	63	45	19
ILL	cif0064	27	77	63	18	3	0	1	0	0	15
ILL	cif0128	26	78	78	78	77	70	19	1	1	22
ILL	cif0256	27	81	81	81	81	81	73	23	4	23
ILL	cif0512	27	81	81	81	81	81	80	31	1	22
ILL	cif1024	26	78	78	78	75	75	72	37	8	25
ILL	vga0128	27	81	78	52	7	0	0	0	0	17
ILL	vga0256	26	78	77	77	73	65	31	8	0	21
ILL	vga0512	26	78	78	78	78	78	75	41	9	24
ILL	vga1024	27	81	81	81	80	81	78	62	27	26
ILL	vga2048	26	78	78	78	78	78	75	60	28	23
ILR	cif0064	27	73	56	15	3	1	0	0	0	4
ILR	cif0128	28	83	84	79	79	67	18	1	0	15
ILR	cif0256	27	80	81	80	81	81	65	17	0	14
ILR	cif0512	26	78	78	78	78	78	74	31	1	23

SG	HRC	Times shown	Row1	Row2	Row3	Row4	Row5	Row6	Row7	Row8	Obj Rec
ILR	cif1024	28	84	84	84	84	84	80	53	2	24
ILR	vga0128	27	80	77	50	6	0	1	0	0	12
ILR	vga0256	27	81	80	75	70	55	21	3	1	16
ILR	vga0512	26	78	78	78	78	78	70	31	9	18
ILR	vga1024	26	78	78	78	78	76	73	52	17	20
ILR	vga2048	27	81	81	80	81	81	75	61	17	25
ILS	cif0064	28	79	82	67	46	20	7	0	0	13
ILS	cif0128	27	81	80	81	81	78	66	30	4	20
ILS	cif0256	27	81	81	79	81	80	81	69	48	24
ILS	cif0512	27	81	81	81	81	81	78	73	62	24
ILS	cif1024	27	81	79	81	81	81	81	72	65	23
ILS	vga0128	27	80	78	64	59	37	25	16	2	15
ILS	vga0256	27	81	81	81	81	75	68	51	34	20
ILS	vga0512	28	84	84	84	84	83	84	80	77	22
ILS	vga1024	26	78	78	78	78	78	78	78	75	24
ILS	vga2048	26	78	78	78	77	78	77	78	74	23
OCL	cif0064	27	81	78	81	81	77	74	56	3	19
OCL	cif0128	28	84	83	84	82	84	84	82	32	28
OCL	cif0256	26	78	78	78	78	78	78	75	44	24
OCL	cif0512	27	81	81	81	81	81	81	79	49	24
OCL	cif1024	26	78	78	77	78	78	78	74	49	25
OCL	vga0128	26	78	78	78	78	75	75	71	41	24
OCL	vga0256	27	81	81	81	81	81	81	81	78	27
OCL	vga0512	27	81	81	81	81	81	81	81	79	24
OCL	vga1024	27	81	81	81	81	81	81	81	81	27
OCL	vga2048	25	73	75	75	75	75	75	75	72	25
OCR	cif0064	27	81	81	80	81	81	77	60	3	8
OCR	cif0128	28	84	84	84	84	84	84	80	22	16
OCR	cif0256	28	84	84	84	83	84	84	81	39	26
OCR	cif0512	28	83	83	82	84	80	82	80	41	26
OCR	cif1024	27	80	81	81	81	80	80	80	43	24
OCR	vga0128	27	81	81	79	81	81	81	79	36	19
OCR	vga0256	27	81	80	81	81	81	81	79	73	25

SG	HRC	Times shown	Row1	Row2	Row3	Row4	Row5	Row6	Row7	Row8	Obj Rec
OCR	vga0512	27	81	81	81	80	81	81	79	77	24
OCR	vga1024	25	75	75	75	73	75	75	74	75	24
OCR	vga2048	27	81	81	81	81	81	81	80	73	25
OCS	cif0064	27	81	81	81	81	80	80	69	51	20
OCS	cif0128	26	78	78	78	78	78	78	73	55	26
OCS	cif0256	26	76	78	78	78	78	78	70	57	24
OCS	cif0512	28	84	84	84	84	84	84	77	62	28
OCS	cif1024	26	78	78	78	78	77	78	72	67	22
OCS	vga0128	27	81	81	81	81	81	81	80	80	25
OCS	vga0256	27	81	81	80	81	81	81	81	81	27
OCS	vga0512	26	77	78	78	78	77	78	78	78	25
OCS	vga1024	28	84	84	84	84	84	83	84	83	28
OCS	vga2048	28	83	84	84	84	84	84	82	82	28
OFL	cif0064	27	81	81	81	80	80	81	73	15	6
OFL	cif0128	28	84	84	83	84	81	82	74	56	17
OFL	cif0256	26	78	78	77	78	78	78	75	61	18
OFL	cif0512	27	81	80	80	80	81	81	77	64	22
OFL	cif1024	27	81	81	81	81	81	81	78	64	21
OFL	vga0128	26	78	77	77	78	77	74	74	65	15
OFL	vga0256	28	83	78	80	81	81	81	82	79	21
OFL	vga0512	28	84	84	82	84	84	84	84	83	26
OFL	vga1024	28	83	84	82	84	84	83	83	82	26
OFL	vga2048	28	84	84	84	84	84	84	84	84	20
OFR	cif0064	26	78	78	77	78	78	78	74	18	7
OFR	cif0128	26	78	78	78	78	78	78	74	42	12
OFR	cif0256	27	81	81	81	81	81	81	81	50	15
OFR	cif0512	26	77	78	78	78	78	78	74	62	10
OFR	cif1024	27	81	81	80	81	81	81	77	64	16
OFR	vga0128	28	84	83	83	84	84	83	82	81	17
OFR	vga0256	28	84	84	84	84	83	84	84	84	17
OFR	vga0512	26	78	78	78	78	78	78	78	77	18
OFR	vga1024	28	83	83	81	79	81	83	84	82	23
OFR	vga2048	27	80	81	81	81	81	81	80	77	20

SG	HRC	Times shown	Row1	Row2	Row3	Row4	Row5	Row6	Row7	Row8	Obj Rec
OFS	cif0064	28	84	84	84	84	81	81	74	55	15
OFS	cif0128	27	81	81	81	81	81	81	75	55	20
OFS	cif0256	27	81	81	81	81	81	80	76	53	22
OFS	cif0512	28	84	84	83	82	81	81	71	51	24
OFS	cif1024	28	84	84	82	84	84	84	75	56	20
OFS	vga0128	26	76	78	78	78	78	77	76	75	21
OFS	vga0256	26	78	78	78	78	76	78	78	78	25
OFS	vga0512	26	78	78	78	78	78	78	77	78	26
OFS	vga1024	28	84	84	82	84	84	84	84	84	23
OFS	vga2048	28	84	84	84	84	84	84	83	81	27

Intentionally Blank

Appendix F PS9 Results by HRC

Table 7. PS9 Result Data by HRC

HRC	Times shown	Row1	Row2	Row3	Row4	Row5	Row6	Row7	Row8	Gen. Elem.	Gender	Hat	Glasses	Necklace	Char.	Pos. ID
cif0064	212	636	636	636	632	610	547	395	89	192	211	205	185	180	147	174
cif0128	212	631	634	636	635	631	627	547	247	193	210	211	199	170	159	178
cif0256	213	638	634	636	639	638	631	590	315	192	212	213	201	177	166	195
cif0512	215	642	644	643	639	641	641	607	364	205	211	214	208	189	182	200
cif1024	216	647	648	648	645	641	642	603	370	196	215	215	209	186	180	201
vga0128	215	641	642	644	645	643	643	610	480	199	215	214	203	189	176	188
vga0256	213	636	638	638	639	637	639	630	604	198	213	209	201	184	171	196
vga0512	214	639	642	641	642	640	641	632	611	192	212	212	210	198	193	190
vga1024	213	639	638	638	638	635	634	634	615	187	210	210	202	192	181	201
vga2048	214	637	642	640	642	640	642	641	616	191	211	212	202	191	182	201

Intentionally Blank