Comparison of interpretability of 1080p30 and 720p60 motion imagery: trade-offs in resolution and frame rate

Charles Fenimore, John Roberts, Hassan Sahibzada (Motion Imagery Quality Metrology Lab)*, and James Yen (Statistical Engineering Division), NIST

Abstract

We have conducted an evaluation comparing the interpretability potential of two standardized HD formats, 1080p30 and 720p60. Despite the lack of an existing motion image (MI) quality scale akin to the NIIRS scale, we have exploited previous work on MI scale development in measuring critical imagery parameters affecting interpretability. We developed a collection of MI clips that covers a wide range of imagery parameters. These well-characterized clips provide the basis for relating perceived image interpretability to motion image parameters, including resolution (related to ground sample distance) and frame rate, and to target parameters such as motion and scene complexity. This report presents key findings about the impact of resolution and frame rate on interpretability. Neither format is uniformly preferred, but the analysis quantifies the interpretability difference between the formats and finds there are significant effects of target motion and target size on the format preferences of the image analysts. The findings have implications for sensor system design, systems architecture, and mission planning.

Keywords: Motion imagery, high definition video, format comparison, 720p60, 1080p30, interpretability

1. INTRODUCTION

The exploitation of motion imagery for interpretation is typically limited by storage and bandwidth. This means tradeoffs must be made among factors such as resolution and frame rate. The optimal tradeoff may be collection dependent, varying with factors including target motion and contrast and on the purpose of the collection. For these reasons, it is important to know how the conditions of a collection affect interpretability. The present comparison of high definition (HD) formats assesses the relative interpretability of two standard HD motion imagery (MI) formats for a limited, but representative, set of collection conditions. The focus is on the international-standard high definition formats 1080p30 (1920x1080 progressive scan at 30 frames per second - Fps) and 720p60 (1280x720 progressive scan at 60 Fps).

This study extends previous work on motion imagery quality scale development to the higher frame rates of the involved formats. Previous evaluations of motion image interpretability measured the effects of collection and target parameters [1]. One particularly relevant study examined the effects of frame rate, target motion, and task complexity [2]. For frame rates in the range from 0 Fps (still imagery) to 30 Fps, there are robust effects of frame rate, motion complexity, and task complexity on interpretability (task completion.) Extending these findings to higher frame rates, with a wider variety of tasks, should provide guidance in mission planning and system design for high performance systems.

The selection of MI materials and interpretation tasks must exercise a significant range of the variation found in applications (Section 2.) The evaluation process (Section 3) involves the comparison of several image formats, including materials originally sourced as 720p60 materials. When down sampled spatially to 480p60 and temporally to 720p30, we term these formats *smallworld*, because it provides a linearly downscaled (2:3) version of the HD format comparison.

In measuring the relative interpretability of 1080p30 and 720p60 video, there is not a single winner for all conditions and targets. Some specific tasks are particularly demanding of resolution, while some targets with particularly high motion call for a higher frame rate. Analysis (Section 4) of the interpretability of the two formats finds a statistically significant dependency on target size and target motion, with 1080p30 favored for smaller targets and 720p60 favored for faster moving targets. Our conclusions and lessons learned are presented in Section 5.

* National Institute of Standards and Technology, Gaithersburg MD 20899-8940; email fenimore@nist.gov.

2. MATERIALS SELECTION AND CHARACTERIZATION

Comparing the relative interpretability of MI clips involves the assessment by imagery analysts (IAs) of a prepared set of clips, having well-characterized attributes. The format comparisons may depend on the methods for obtaining the materials and on their selection, preparation, and presentation as well as the design of the test instrument used to obtain information from the IAs. This section describes our approach and methods used in developing the MI materials and the test instrument. The design was guided by the requirement to develop combinations of clips and interpretation tasks for which changes in frame rate and resolution might make a difference.

2.1 Motion Image (MI) materials source and preparation

The study was constrained by the limited availability of suitable source material in the 1080p30 and 720p30 formats. Since changes in the smallest details of scene or motion affect the limits of interpretability, it is not sufficient to select different clips in the two formats, regard them as "comparable" and compare their interpretability scores. It is important to compare clips in the two formats that are as nearly identical as possible. Two possible approaches are: (1) simultaneous ("boresighted") shooting of a scene in the two formats, using two cameras, or (2) derivation of matching 1080p30 and 720p60 clips from a 1080p60 original (the 1080p30 by discarding every other frame, and the 720p60 by scaling the 1080p frames down to the 720p format). Neither approach was available for this study.

The study team had available a substantial amount of content shot in 1080p30 and 720p60 (but not simultaneously) at the 2003 Marine Corps Marathon and at the airfield at Patuxent River Naval Air Station. We also had access to additional content captured in the 720p format. The solution found to produce clips pairs for comparison was to play the 1080p30 content at double speed, to produce nominal 1080p60 content, from which matching 1080p30 and 720p60 clips were derived as described above. The tradeoff was that the action in the 1080p60 clip and the derived 1080p30 and 720p60 clips is twice normal speed – this is an unavoidable part of the derivation process. This means that a particular action in the original (normal speed) clip that may have been an interesting target for an evaluation task (for example a running person) is now too fast in the derived clips for the same sort of tasks. However, since the original 1080p30 content available to the MIQM team is rich in types and speeds of motion, there are actions that were too slow to be of interest in the original (for example a person walking slowly) that now become fast enough to support a useful interpretation task in the speeded-up clips. It may be useful to visualize the derived clips as accurate and internally consistent depictions of a hypothetical *fastworld*, where everything happens at double speed, but the video collection and playback are consistent with 1080p30 and 720p60, and only the selection of tasks differs from the real-time playback. The results of the study may be used to determine whether *fastworld* fully supports the interpretability comparison of 1080p30 vs 720p60.

Prior to the evaluation, it was suggested that the double speed *fastworld* clips be played at half speed, so that the speed of motion matches that of real-time playback. This form of viewing is effectively a comparison of 1080p15 and 720p30, and does not address the highest frame rates of interest in this study.

In order to build a conveniently transportable evaluation platform, we performed MPEG-2 video compression at 20 Mb/s on the imagery. This permitted us to present the evaluation on a PC, without the significant loss of quality that would have resulted from compression at lower bit rates.

In order to assess the impact of the spatial format conversions and video compression on interpretability, detailed interpretation tasks were performed on selected matching frames in the 1080p and 720p formats prior to the test. Informal comparison found the interpretability of the 1080p frames when scaled and compressed to 720p were identical to that of the comparable clip from the 720p source. The comparison provides confidence in the validity of the content selection and test preparation as far as spatial conversions and video compression are concerned.

In addition to the principal "*fastworld*" comparison of 1080p30 to 720p60, we developed "*framerate*" content to evaluate high-rate frame rate (FR) effects on interpretability, with paired 720p60 and 720p30 clips. Our prior work on FR[1] found no statistically significant difference in interpretability when comparing motion imagery presented at 30 and 15 Fps over a range of motions and targets. A similar indifference between 30 and 60 Fps, over the wider range of motions used in this study, would be strong indication of a preference for 1080p30 for similar targets with similar motion characteristics. Because we used original 720p60 materials, there were no possible confounding results from *fastworld* effects.

Finally, an additional derivation from the available 720p60 clips was added to the evaluation as a further check on the tradeoff between resolution and frame rate. In analogy to the derivation of matching 1080p30 and 720p60 clips from an (unavailable) 1080p60 original, matching 720p30 and 853x480p60 clips can be derived from the (available) 720p60 originals. The production involves scaling and down sampling, such as was applied to the *fastworld* 1080p60, but since the originals are 60 Fps, there is no need to double the frame rate. Because the nominal linear resolution is 2/3 that of the core formats, 1080p30 and 720p60, we term these materials *smallworld*. Even though the resolution is lower for the *smallworld* clip pairs than for the target formats, the frame rates are the same, the relative tradeoff between resolution and frame rate is the same, and the speed of target motion (relative to frame size) is the same as in the original clips. Thus *smallworld* may provide independent confirmation of effects seen in *fastworld*. In addition, unlike the available 1080p30 materials, some of the original 720p60 was collected from aerial platforms, an important image source for reconnaissance.

2.2 Selection of evaluation content: clips and interpretation tasks

Previous work on MI quality has found interpretability depends on the motion of the target and the complexity of the scene, in addition to factors found in still image quality characterization such as ground sampled distance (GSD, a measure of nominal resolution), contrast, edge response, and noise. Guided by these findings, we aim to select clips with a wide range of motion characteristics, from very fast or complex motion down to scenes with very simple motion and with little or no motion. As for resolution, the selection of identical scenes for the comparisons means that the ratio in GSD is fixed at 3:2 in every clip pair. Nevertheless, we expect resolution-dependent differences would be most significant near (but above) the sampling theorem limit of 2 pixels per target feature length [3].

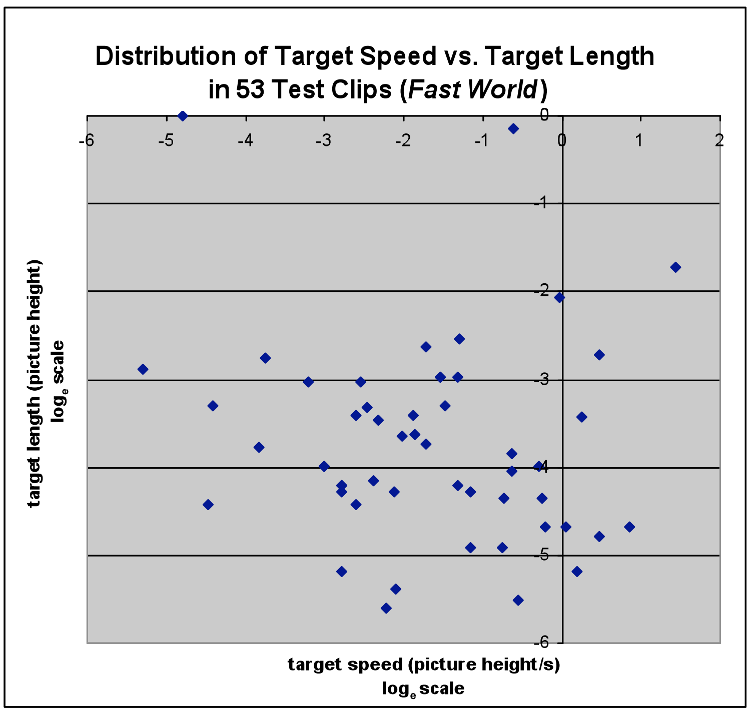

Figure 1: Distribution of Target Speed and Target Length for 57 Test Clips: The distribution provides coverage of the region with higher target speeds (above 0.05 *H/s*, where *H* is in units of picture height) and with smaller targets of size < .05**H*. The authors expect the most significant resolution-related effects to arise on smaller targets; note the scarcity of targets between 2 and 8 pixels in length. Frame rate effects are expected to be most significant for high-speed targets.

In support of clip selection and analysis of the IA responses, the target motion and resolution characteristics of each clip are measured prior to the evaluation. The speed of motion is an objective measure while motion complexity involves rater judgment. For this study we have measured target length in units of picture height (*H*) and the speed of motion as target motion relative to the immediately surrounding background in units of *H per second (H/s)*. These measures of length and speed do not depend on the format in which the measurement is made.

Figures 1 and 2 are scatter plots showing the measured values of target lengths and target speeds for the *fastworld* and the *smallworld* portions of the evaluation. The range of values suggests the clip and task selections are sufficiently well distributed for the measurement of target size and motion effects. Whether the number of samples supports high confidence in those measurements is addressed in Section 4 (Analysis.)

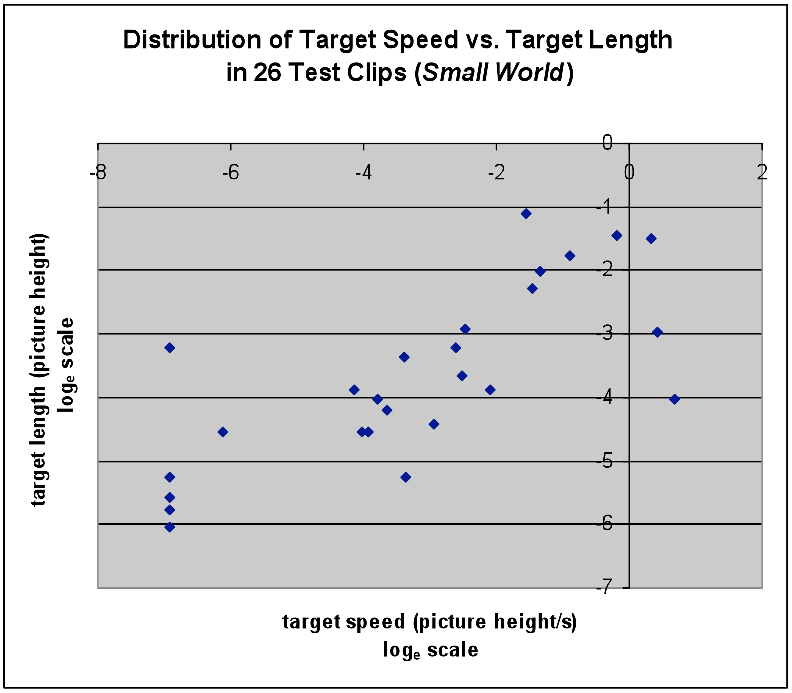

Figure 2: Distribution of Target Speed and Target Length for 26 Test Clips in the *smallworld* clips: The covered region is not as large as for Figure 1. The lower degree of coverage and the smaller number of samples may impact statistical confidence in the expected effects due to motion. Targets with length between 2 and 8 pixels are not found in *fastworld*.

Clips and tasks are selected to encompass a wide range of motion, to maximize the chances of finding any crossover between preference for high resolution and preference for high frame rate. Clips and tasks are also selected to provide a good fill across the most "useful" range of motions, matching when possible the MIQM motion imagery criteria being developed through the National Geospatial-Intelligence Agency. Tasks range from high motion, such as identification and tracking of fast human motion or identification of features on fast moving vehicles, to identification of features on motionless or slow moving people, vehicles, and background objects. The measured speed for "fastworld" clips is based on the motion in the speed-doubled version.

The following information is recorded for each clip used in the evaluation, to aid in analysis of the results:

- Target motion: Measured speed recorded in multiples of frame height per second.
- Target motion period, in units of seconds (s).
- Motion complexity. Subjectively binned High, Medium, and Low. As a reference, vehicular motion would have Low complexity, and most human motion would have High complexity.
- Target length in the direction of motion, measured in units of frame height (*H*) in the direction of target motion.
- Target GSD (ground sampled distance). The estimated lateral distance on the target between two image pixels.
- Minimum resolvable target feature size (in cm): If the image is clear and in focus, the size is that of one pixel, otherwise it is the smallest feature size on the target that can be resolved.
- Contrast. Subjectively binned High, Medium, and Low.

2.3 Playback method and display modes

The tests are performed on a desktop computer configured for 1920x1200 display, using a 23" LCD monitor. The open-source VLC player[1] is used for clip playback. 1080p30 and 720p60 clips for comparison showed the same scene (matching features are the same percentage of frame height in the two formats). The ambient lighting and screen brightness conform to the ITU standard for evaluation of television image quality (BT500-12) [4]. Viewing distances are self-selected by the IAs.

1080p30 clips are played in "actual pixels" format (a pixel on the screen represents exactly one pixel in the image). For playback of 720p60, we compare two playback modes: "actual pixels" (so the image frame size is less than that of the 1080p30 playback), and "full screen" (for which the 720p image frame would be the same size as the 1080p30 image frame). Inspection of sample clips playing in the two display modes found there were negligible differences in apparent interpretability between "full screen" and "actual pixels". In view of these findings, we play all our clips (including 720p, 1080p, and 480p) in the "actual pixels" format. A small ancillary evaluation makes a small number of direct comparisons of interpretability for the same 720p60 clip in "full screen" and "actual pixels" playback modes.

To minimize inter-rater variation, the playback format is fixed; for example, the IAs view each clip a specific number of times. Also, prior to viewing each clip, we provide the IAs with suitable visual and written cues whenever the length of the clip or other factors might make task assessment difficult.

3. EVALUATION PREPARATION AND EXECUTION

The evaluation is divided into two parts, based on the way the analyst responds to the questions. For each task in Part 1, the IA views a single clip (single stimulus) and provides a rating based on that clip. Part 2 is double stimulus, asking the analyst to perform pairwise comparisons of the two formats. Both Part 1 and Part 2 contain elements to support the *fastworld*, *framerate*, and *smallworld* tests.

3.1 Evaluation control and scoring

A program running on a personal computer controls the evaluation, invoking image display, presentation of questions, and capture of the IA's responses. Pacing of the questions is under IA control. The playback of each clip is fixed and is not subject to IA control. Instructions and a short training precede the evaluation. Each task is presented to the IA by listing the task in text form and showing a "cue image" indicating the region of interest for the task target. When the IA is ready, the program shows the clip(s) (one clip for Part 1 of the test, or two clips, shown in "A/B/A/B" sequence for part 2). The program presents the task questions a second time and prompts the IA to enter the scores:

The IA scores each question on either a 5-point or a 3-point scale. Part 1 of the test presents the IA with two questions:

- First Question: "Using the clip, rate your confidence in your ability to . . . "
  Scoring scale: 1 = "no confidence", to 5 = "fully confident"
- Second Question: "How easy would it be to use this clip in accomplishing this task?"
  Scoring scale: 1 = "very difficult to complete the task with this clip", 2 = "some difficulty", 3 = "easy to complete".

Part 2 of the test (pairwise comparison) also presents the IA with two questions:

- First Question: "Which clip gives you greater confidence in your ability to…?"
  - Scoring scale: 1 = "Clip A provides much greater confidence than Clip B", to 5 = "Clip B provides much greater confidence than Clip A"
- Second Question: "Which clip is more useful in allowing you to…?"
  - Scoring scale: 1 = "Clip A is much more useful than Clip B", to 5 = "Clip B is much more useful than Clip A".

---

[1] The mention of products is not a product endorsement by NIST, nor is it meant to imply the product is the best available for the purposes of this study.

Analysts are given a questionnaire afterwards, soliciting their impressions of the questions asked, how the test went generally, and how it could have been improved. These responses are used to evaluate the test's results and reliability.

4. ANALYSIS

The analysis of our data is described in two stages. The basics of treatment of the data and measures of data consistency, as well as some descriptive statistics relating the several components of the study are presented in 4.1. The analysis of specific imagery attributes on the quality differences between the candidate formats is in 4.2.

4.1 Data consistency and analysis

4.1.1 Canonical scale of the data:

Since the answers in Part 2 are symmetric in their wording with respect to format (Clips A and B), the data from *fastworld* and *smallworld* was transformed so that all answers follow a canonical 1-5 scale in which higher values favor FR=60. Thus a score of 5 indicates that the clip with FR=60 provided much greater confidence or was much more useful than the clip with FR=30; a score of 1 indicates the opposite strong preference in favor of the FR=30 clip, while a score of 3 indicates no preference between the two clips.

4.1.2 Relative ease and relative confidence:

For *fastworld* and *smallworld*, each question contains both a relative confidence and a relative ease of use component. The responses are highly correlated; for *fastworld*, relative ease and relative confidence have a correlation 0.88; for *smallworld* the correlation is 0.82. For *fastworld* there were no cases where an analyst found one clip easier to use but the other clip inspiring more confidence; for *smallworld*, there was only one such anomalous case.

Figure 3 shows the average relative ease scores for each question plotted against the average relative confidence scores for that question. The averages from Part 2, depicted by circles, are very well associated. The Part 1 points, depicted by triangles, follow a different association because the second component of the Part 1 questions are on a 1-3 scale rather than a 1-5 scale.

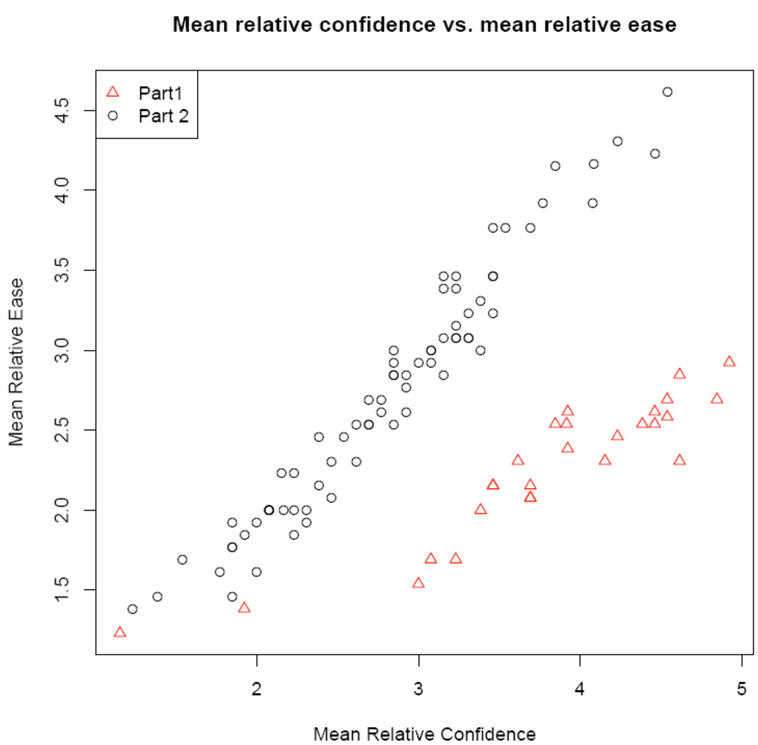

Figure 3: Each point plots the average relative ease for a question versus its average relative confidence (Part 2, triangles) or its ease vs. its confidence (Part 1, circles).

The evaluation engaged 13 image analysts. One indication of the degree of consistency between the analysts is the correlation between the scores of the analysts. For a given pool of questions, there are (13x12)/2=78 correlations between the 13 analysts. Using the scores from *fastworld* and *smallworld*, Figure 4 shows a histogram of these correlations between analysts for the relative confidence scores. All the correlations are positive and between 0.1 and 0.7.A similar analysis of the ease-of-use data, finds one anomalous pair of analysts having a correlation slightly below zero on the *fastworld* and *smallworld* relative ease questions. All the other correlations are positive.

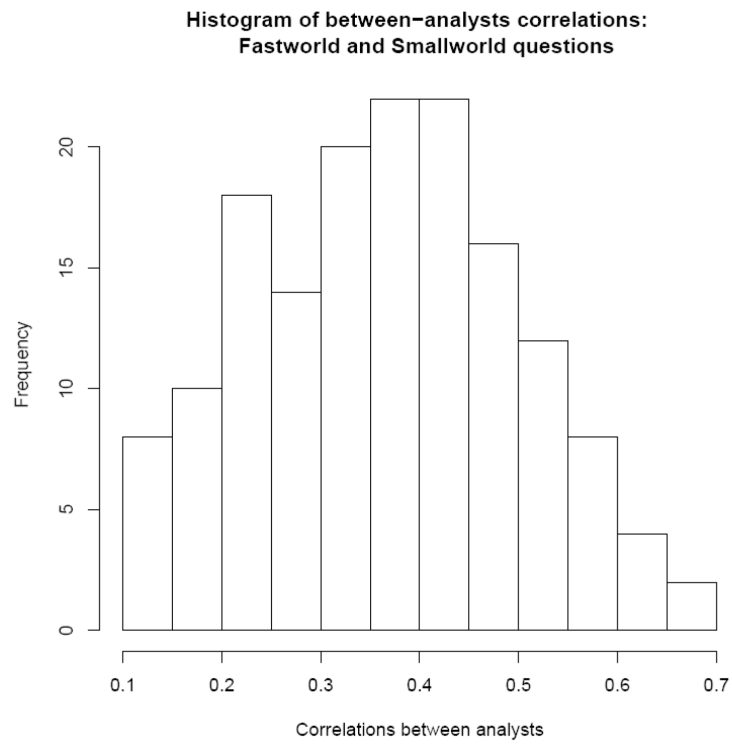

Figure 4: The between-analyst correlations for the relative confidence for the *fastworld* and *smallworld* questions. All the correlations are positive and between 0.1 and 0.7.

That the *fastworld* and *smallworld* data results are restricted to the values 1, 2, 3, 4, and 5 makes the correlation a somewhat crude tool. An alternative view of the data examines the explicit agreements and disagreements between analysts in each question.

For any question in *fastworld* or *smallworld*, we consider two analysts to have a total disagreement if they prefer different clips; they agree if they prefer the same clip or both have no preference between the clips; and they have a marginal disagreement if one analyst has no preference, whereas the other analyst does prefer one of the clips. We look at all pairs of analyst responses from each of the 55 relative confidence questions in *fastworld* and *smallworld*. About 10 percent of the response-pairs had total disagreements (they preferred different clips), 46 percent had total agreement (they preferred the same clips or both had no preference between the clips), and 44 percent had partial disagreement (only one analyst has no preference). As with the correlations, this indicates that a trend towards agreement between the analysts coupled with considerable variability between the analysts.

4.2 Factors from imagery attributes

4.2.1 Factors - Clip parameters

In modeling the dependencies of the *fastworld* data on single factors for which numerical data is available, the target speed is the only such factor that is highly significant. Figure 5a depicts that data for the *fastworld* data of Part 2. Note that the range of the numerical factors makes a logarithmic transformation appropriate (log base *2* is used); however, since zero is often included among the factor values, a small increment, denoted by *incr* is added to the data before the log transformation to avoid infinite or undefined values (if the factor values are all positive, then *incr*=0; otherwise, *incr* is one-third of the smallest non-zero value). Also, since the raw data takes the form of the integer values 1-5, standard scatter plots of the raw data would not be useful. Instead, the mean relative confidence, averaged over the 13 analysts, is plotted for each question in the section. Figure 5b shows the analogous plot for relative average ease, which is very similar to the relative confidence plot. Both plots show an upward trend coupled with considerable noise.

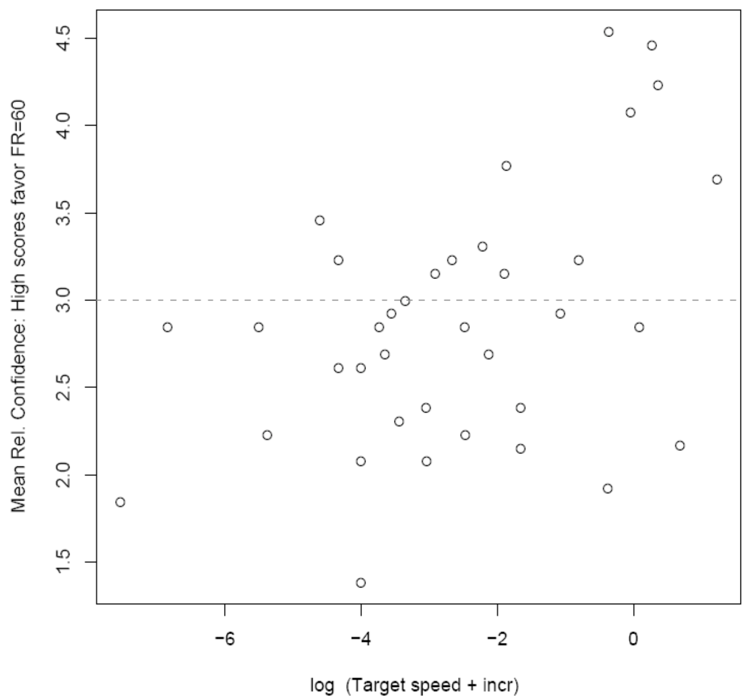

Figure 5a: A statistically significant effect of target speed is found in the measured confidence in *fastworld*. The strongest preference is found for values of the speed close to 1.0 frame heights /s.

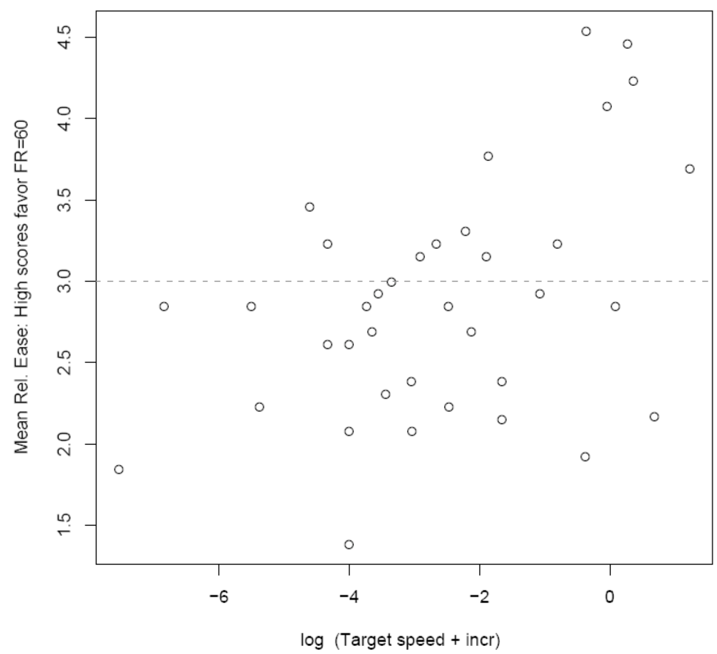

Figure 5b: Similarly to Figure 5a, the target speed has a significant effect on the measured *ease of use* in *fastworld*. The strongest preference is found for values of the speed close to 1.0 frame heights / s.

Figure 5c shows a similar trend for the Part 1 confidence data. The vertical coordinate of each point depicts for a certain question in Part 1 the difference between the average confidence for those analysts who saw the clip in 720p60 and average confidence of those who saw in 1080p30. The upward trend shows that increased speed of the target favors the 720p60 format, even as the majority of questions produce a preference for 1080p30.

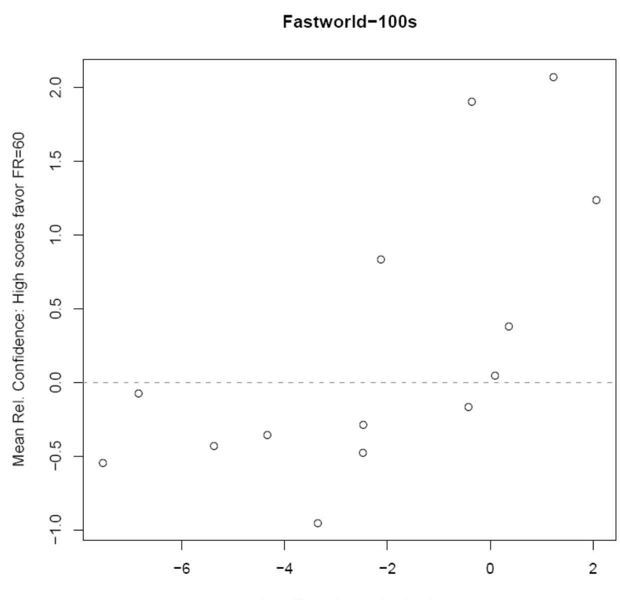

Figure 5c: The vertical coordinate of each point depicts for a certain question in Part 1 the difference between the average confidence for those analysts who saw the clip in 720p60 and average confidence of those who saw in 1080p30. The upward trend shows that increased speed of the target favors the 720p60 format, even as the majority of questions produce a preference for 1080p30.

However, Figure 5d, which is the analogous relative confidence plot for the *smallworld* data, shows a much more clear-cut upward pattern. This difference may be attributable in part to the differences between the *smallworld* and *fastworld* factor coverages seen in Figures 1 and 2 that highlights possible association between target length and target speed. In the *smallworld* data, there seems to be an association between target length and target speed for much of the data.

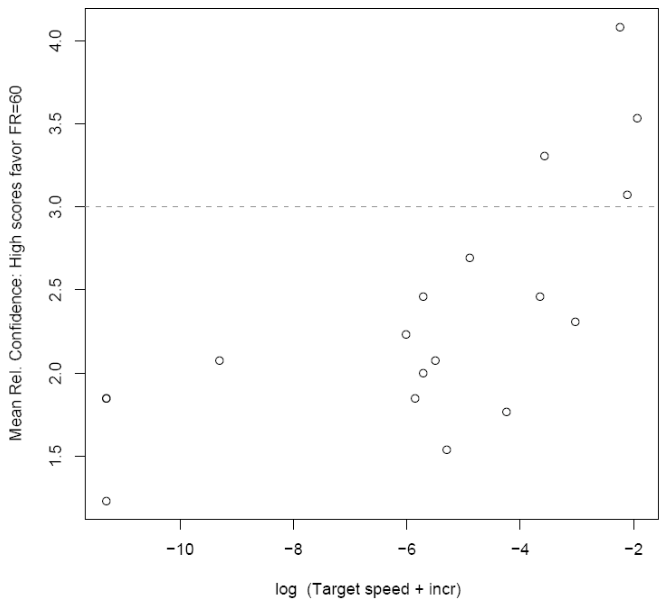

Figure 5d: In *smallworld* the association between target speed and relative confidence is even clearer than in *fastworld* (5a-5c)

Figures 5e and 5f plot the mean relative confidence versus log target length for *fastworld* and *smallworld*, respectively. There is a significant target length effect in *smallworld* but none for *fastworld*. This may result from the small, resolution-limited size of targets in *smallworld*, seen in Figure 2.

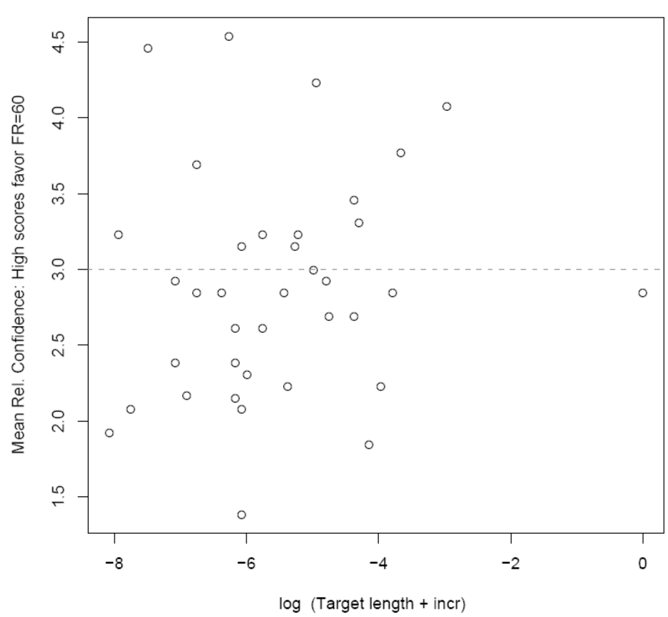

Figure 5e: For the *fastworld* data, there is no clear trend between target length and mean relative confidence, although some parts of the data may seem suggestive.

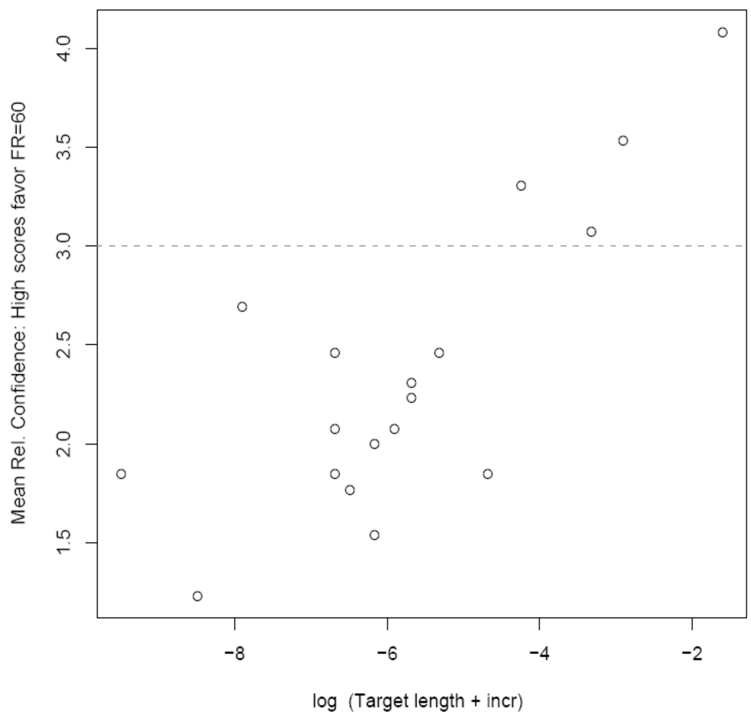

Figure 5f: Unlike the *fastworld* data, the *smallworld* data shows a strong association between target length and mean relative confidence (here, smaller target length are associated with a greater preference for more resolution). This may also be an indication of association with target speed. Unlike *fastworld*, the *smallworld* data has several very small targets. The different MI clip selection may permit seeing the interpretability differences.

4.2.2 Modeling Confidence and Ease of Use in *fastworld* and *smallworld*

A possible prediction model involves regression of relative confidence scores, C_R, in *fastworld* on the log (base 2) of both the target length, L_t, and the target speed, S_t. The fitted model is:

$$C_R = 4.8 + 0.25 * log(S_t + \text{í}_S) + 0.23 * log(L_t + \text{í}_L), \quad (1)$$

where í_S and í_L are small increments added to the variables to ensure positivity before taking the log transform. As described in previous sections, í is one-third of the smallest positive value present, unless all values are already positive in which case í=0. The estimated coefficients with their estimated standard errors are 4.8 (0.3), 0.25 (0.03), and 0.23 (0.4). All the terms in the regression are statistically highly significant (p-values less than 0.0001); however, the model only accounts for about 13 percent of the variability present. An interaction term for $log (S_t + \text{í}_S)$ and $log(L_t + \text{í}_L)$ was found not to be significant. Note that the intercept 4.8 is large because the input variables, being log transformed, are negative for most of the sample space. The model indicates that greater target speed increases preference for 720p60, and small target length increases preference for 1080p30. Note that unlike in figures in Section 4.2.1, this model is calculated using the individual relative confidence scores rather than the question averages.

An analogous fitted model for *smallworld* is

$$C_R = 4.9 + 0.27 * log(S_t + \text{í}_S) + 0.38 * log(L_t + \text{í}_L) + 0.03 * log (S_t + \text{í}_S) * log(L_t + \text{í}_L). \quad (2)$$

This model is calculated using the individual relative confidence scores. The estimated coefficients with their estimated standard errors are 4.9 (0.4), 0.27 (0.07), 0.38 (0.7), and 0.03 (0.01). All the terms in this regression are statistically significant (p-values less than 0.002), including an interaction term for $log\ (S_t + \acute{\imath}_S)$ and $log(L_t + \acute{\imath}_L)$. As in *fastworld*, the estimated model again indicates that greater target speed increases preference for the faster frame rate, and small target length increases preference for the greater number of pixels. The estimated model accounts for 27 percent of the variability present, which while more than the *fastworld* model, still indicates considerable variability in the responses. The model is similar to the *fastworld*-estimated model; the differences may be due to the different formats used in *smallworld*, and also the possible correlation between target speed and target length seen in the *smallworld* data clips.

5. CONCLUSIONS

Measured over the entire collection of pairs of MI clips and interpretation tasks used in this study, there is a slight but statistically significant preference for the higher spatial resolution / lower frame rate format (1080p30) over the lower resolution / higher frame rate (720p60). This overall average preference is overlaid with (1) a marked preference for 720p60 for very high-speed targets, typically traveling at speeds ranging from 0.1 to 1.0 frame-height/sec; and (2) in one portion of the study, a preference for the higher resolution format (1080p30) for small targets having lengths of 2 – 8 pixels, slightly above the sampling theorem limit. The latter result is only seen in the *smallworld* part of the study, which appears to be a consequence of the lack in *fastworld* of targets near the resolution limits (Figure 1) implied by the sampling theorem applied to *spatial* frequency. For complex tasks, more resolution may be required. The modulation transfer function (MTF) of the imaging system provides a more reliable estimate than does the sampling theorem of the resolution required for interpretability. The highest speeds and the smallest targets are only achieved for a small number of the clip/task pairs in our study. In MI applications for which the range of image target sizes and speeds is known, it may be possible to establish a preference for one or the other of the two formats.

The application of our results to other formats is not entirely clear. In particular, extrapolation of this methodology to very low frame rates will require testing at the frame rates in question. The previous work on frame rate effects[1] found the impact of frame rate was task dependent and large at low frame rates. On the other hand, our results on the preference for the higher resolution format for small targets, certainly comes as no surprise. In particular, for applications falling in the asymptotic case of very low target motion the preference for more resolution may be regarded as a result of the still imagery NIIRS.

Finally, it is worth noting the results of the study confirm the validity of the use of *fastworld* and *smallworld* to enable this interpretability study. Nevertheless, in planning for future studies, the capture of new content collections in either bore sighted (simultaneous capture) 1080p30 and 720p60 or in 1080p60 format would eliminate the need for this device. While the results of such studies would presumably be comparable to those of this study, the demands on the analyst would be diminished.

6. ACKNOWLEDGEMENTS

We wish to acknowledge the assistance of the National Geospatial-Intelligence Agency in this work. To our knowledge, this paper is the first published work employing the *Marine Corps Marathon* and *PaxRiver* clips. The team that collected and provided these materials includes Bill Butterfield, Henry Dardy, Gary Demos, Ronald Lee, and George Palmer. We also acknowledge the contributions of Fred Petitti and Rachel Bowers, who assisted in developing the design.

7. REFERENCES

[1] Fenimore, C., Irvine, J., Cannon, D., Roberts, J. et al., "Perceptual Study of Frame Rate Effects in Motion Imagery Interpretation", Proc SPIE 6057, 60570P-1:9 (2005).
[2] Irvine, J., Aviles, A.I., Fenimore, C., Roberts, J., et al., "Motion Image Interpretability Scale Development", Optical Eng. Vol. 46(11), 117401-1:12 (2007).
[3] Lim, J.S., Two-Dimensional Signal and Image Processing, Prentice-Hall, New York, (1990).
[4] International Telecommunications Union, "Methodology for the subjective assessment of the quality of television pictures", Recommendation BT500-10, (2000).