

Reverse Conducting Evaluation

Craig Stuart Sapp <craig@ccrma.stanford.edu>

15 Nov 2005

This is a generic *Mathematica* notebook used for examining the individual reverse conducting performances of Chopin Mazurkas compared to carefully corrected versions of the average beat tappings for the performance.

Mazurka in A Minor, Op. 7, No. 2 Friedman 1930

■ Loading the raw data

First load the absolute beat positions for the average beat positions and the corrected beat positions:

```
performance = "pid5667230-09";
pvar = StringReplace[performance, "-" -> "x"];

Get[performance <> ".macorrected", Path -> "D:\mazurkas\eval"]
mcbeats = ToExpression[pvar <> "corrected"];

Get[performance <> ".maaverage", Path -> "D:\mazurkas\eval"]
mabeats = ToExpression[pvar <> "average"];
```

Now measure the time difference between each average and corrected beat. A negative value means that the average beat occurs before the corrected beat time.

```
mdiffs = Transpose[mabeats][[2]] - Transpose[mcbeats][[2]]

{246, 169, 75, -42, 42, 54, 18, 174, 0, -3, 57, -4, 83, 79, 12, -16, 64, 0, 41, 99, 68, 49, -5, 8,
241, 103, 21, -79, 66, 60, 34, 241, 83, 2, 21, -26, -91, 22, 31, -11, -51, -43, 51, 205, 50,
-68, -199, 42, 127, 137, 63, 12, 47, 43, -5, 32, 37, 4, 10, 69, -33, -117, -116, 36, 31, 58,
-11, -28, 65, -19, 6, 8, 65, 0, 29, -17, 78, 63, 122, 232, 63, -6, -10, 90, -40, 24, -7, -86,
-63, 53, 20, 124, 39, 0, -71, 70, 91, 132, 70, 0, 66, 87, 88, 92, 76, 47, 56, 95, 64, 23, -123,
19, -28, 0, 0, 41, 29, -6, -39, 0, 34, 66, -44, -89, 77, -28, 0, 202, 79, -29, 27, 17, -32, 14,
7, -18, -15, -126, -105, -263, -48, -64, -10, 48, -76, -161, 61, 39, 30, 50, 54, 81, 78, 17,
55, 60, 15, 41, 50, 12, 91, 41, 23, 101, 80, 31, 132, 52, -16, 128, 40, 47, 89, 31, 24, 68, 74,
41, 129, 19, 53, 50, 54, 67, 84, 47, 48, 74, 60, 79, 105, 67, 122, 169, -151, 135, 72, 67, 43,
67, 91, 12, 52, 75, 44, 62, 63, 27, 89, 56, 18, 90, 65, -15, 88, 122, 12, 37, 21, -15, 15, 27,
12, 19, 48, 40, 106, 24, 19, 11, 0, 0, 45, 62, 36, 76, 74, 0, -26, 0, 0, -194, 69, -16, -24, 53,
59, 30, 44, -22, 92, 103, 32, 4, 0, 6, 80, 72, 76, 50, 51, 11, 135, 118, 270, 43, 6, -47, 119,
83, 17, 194, 24, -9, 25, 0, 125, 101, 119, -66, 51, 0, 27, 81, 93, 71, 0, 78, -73, -220, 59,
-17, 113, 123, 14, 238, 0, 19, 41, 109, 33, 10, 38, -70, -66, 0, 0, 185, -17, -49, -43, 248, 0}
```

The average correction to each beat from the average tap times of all trials is typically around 48 milliseconds. This means that, on the average, each average tap beat position must be moved 48 milliseconds to align with the correct beat time.

```
<< Statistics`ConfidenceIntervals`
```

```
Mean[Abs[mdiffs]] // N
```

```
59.6454
```

■ Manual correction of average reverse conducting beat positions

This section displays the difference between the average reverse conducted beat positions and the values derived from manually correcting these values by ear using a soundfile editor.

This corrected beat duration data was measured by listening to the audio file and manually locating the beat positions in the soundfile. The initial positions of the average reverse conducting performance were used as a baseline, and this baseline was then adjusted so that it sounded as if all beats were occurring with the attacks of notes in the performance. There are probably a few errors in the corrected due to the tediousness of getting this data, and there are a few locations in the audio where the beat positions become vague, but overall, this data represents a highly accurate description of the pianists' actual beat-by-beat performance tempo, probably an accuracy within 10 milliseconds for all beats.

```
corrected = Transpose[mcbeats][[2]]
```


```
{1931, 2483, 2966, 3483, 3851, 4200, 4697, 4957, 5485, 5850, 6146, 6570, 6974, 7346, 7692, 8118, 8465,
 8925, 9329, 9683, 10130, 10533, 11109, 11765, 12183, 12677, 13220, 13750, 14088, 14438, 14892,
 15058, 15531, 15911, 16217, 16639, 17207, 17630, 17939, 18343, 18805, 19330, 19839, 20235,
 20841, 21461, 22224, 22889, 23495, 23892, 24332, 24829, 25276, 25672, 26149, 26585, 26975,
 27441, 27902, 28285, 28833, 29693, 30364, 30819, 31315, 31705, 32195, 32706, 33090, 33620,
 34136, 34654, 35438, 36100, 36782, 37254, 37592, 37988, 38423, 38672, 39182, 39606, 39977,
 40298, 40829, 41250, 41628, 42151, 42730, 43289, 43942, 44489, 45161, 45754, 46452, 47110,
 47701, 48107, 48557, 49084, 49524, 49900, 50339, 50738, 51104, 51544, 51960, 52305, 52755,
 53422, 54055, 54514, 55018, 55454, 55917, 56367, 56788, 57286, 57845, 58400, 59273, 60025,
 60843, 61396, 61722, 62178, 62701, 62973, 63483, 63922, 64309, 64708, 65247, 65693, 66043,
 66474, 66941, 67623, 68365, 69201, 69931, 70547, 71097, 71549, 72190, 72940, 73547, 73964,
 74355, 74752, 75117, 75424, 75738, 76122, 76443, 76746, 77123, 77437, 77746, 78129, 78417,
 78766, 79137, 79413, 79716, 80089, 80368, 80726, 81226, 81604, 82064, 82408, 82714, 83061,
 83434, 83745, 84101, 84478, 84738, 85106, 85607, 86014, 86390, 86745, 87057, 87398, 87752,
 88124, 88728, 89208, 89603, 90218, 90866, 91498, 92532, 92877, 93262, 93627, 94002, 94314,
 94618, 95015, 95328, 95621, 95996, 96308, 96615, 97007, 97297, 97625, 98014, 98292, 98604,
 99049, 99358, 99673, 100264, 100789, 101259, 101677, 102039, 102352, 102721, 103091, 103443,
 103792, 104056, 104412, 104922, 105369, 105762, 106117, 106432, 106721, 107097, 107455,
 108056, 108660, 109235, 110161, 111006, 112101, 112923, 113484, 113968, 114368, 114733,
 115094, 115399, 115818, 116071, 116357, 116746, 117092, 117428, 117785, 118084, 118397,
 118735, 119084, 119389, 119840, 120129, 120544, 121837, 122804, 123505, 124046, 124339,
 124711, 125239, 125510, 126026, 126420, 126813, 127259, 127705, 128099, 128396, 128932,
 129328, 129803, 130224, 130604, 131047, 131461, 132029, 132648, 133429, 134285, 134863,
 135392, 135684, 136017, 136545, 136719, 137302, 137633, 137989, 138315, 138806, 139275,
 139632, 140144, 140677, 141250, 141940, 142520, 143596, 144552, 145720, 146582, 148323}
```

```
avgtimes = Transpose[mabeats][[2]];
```

Here are the timing differences between the average reverse-conducting absolute times and the absolute times measured in a sound editor:

```
diffs = corrected - avgtimes;
```

```
ListPlot[diffs, Frame → True, AspectRatio → 1 / 4, PlotStyle → PointSize[0.008]];
```




The average beats are not quite centered on the corrected beats. The average beats are, on the average, 4.6 milliseconds after the corrected beats.

```
offseterror = Mean[diffs] // N
-35.8307
```

The histogram below shows the data from the plot above. It generally shows a nicely distributed range of corrections. The peak at 0 is due to the correction values being too small to change to the theoretically correct values about up to 10 milliseconds before their measured positions.

```
<< Graphics`Graphics`
```

```
Histogram[diffs, AspectRatio → 1 / 4, Frame → True, BarEdges → False,
HistogramCategories → 50, HistogramScale → 100, HistogramRange → {-200, 200}];
```



■ Trial Quality Measurements

Define a logarithmic score for the amount of correction needed for a given beat. A score of 1 means 20 ms correction, 2 means 40 ms correction, 3 means 80 ms correction, 4, means 160 ms correction, 5 means 360 ms correction, *etc.* A score of 0 means there was less than 1 millisecond correction necessary, and a score of 0.5 means 10 milliseconds were needed to correct the beat location.

Qualitatively, a score of 0-1 is an "A", 1-2 is a "B", 2-3 is a "C", 3-4 is a "D", 4-5 is an "E", and 5-6 is an "F". A score of 6 means the correction was 640 milliseconds, or about 2/3 of a second.

```

QScore[seq_] := Module[ {sign, scores},
  sign = Map[If[# < 0, -1, If[# > 0, 1, 0]] &, seq];
  scores = Log[2.0, Abs[seq] / 10] /. -Infinity -> 0;
  scores = Map[If[# < 1, (# + 1) / 2.0, #] &, scores];
  scores = Map[If[# < 0, 0, #] &, scores];
  scores * sign
]

```


Mostly the corrections for the average of the tapping trials are evenly distributed between A, B, C, and D quality, although "A" corrections were the most common. The worst case correction was between a value of 4 and 5. This occurs on the first beat when there has been no preparation of a previous event, so is most likely due to the reaction time of the reverse conductor. Other corrections in the range between 4 and 5 are also likely due to an unprepared change in tempo which was not expected beforehand.

```
<< Graphics`Graphics`
```

```

Histogram[Abs[QScore[mdiffs]], AspectRatio -> 1 / 4,
HistogramScale -> 100, Frame -> True, HistogramCategories -> 5];

```



The range from 0-1 can be characterized as "no audible difference" (0-20 ms); 1-2 as "slight audible difference" (20-40 ms); 2-3 "minor audible difference" (40-80 ms); 3-4 as "noticeable audible difference" (80-160 ms), 4-5 as "very noticeable audible difference" (160-320 ms), and anything higher as "extremely noticeable audible difference" (> 320 ms).

```
Mean[Abs[QScore[mdiffs]]]
```


```
2.11351
```

Here is a histogram display of the necessary corrections with positive and negative adjustments separated. Negative values mean that the correct value occurs before the average reverse conducting time. Notice that most of the larger corrections occur in the negative range

```

Histogram[QScore[mdiffs], AspectRatio -> 1 / 4,
HistogramScale -> 100, Frame -> True, HistogramCategories -> 11];

```



■ Examine the learning curve

How does the quality of the reverse conducting improve over time? Does the conductor learn to follow the performance better after repeated listening? First, load the individual tapping trials which contain precalculated offset values into the audio file of the performance recording:

```
Get[performance <> "-abs.ma", Path -> "D:\mazurkas\eval"]
data = ToExpression[pvar <> "abs"];
```

Now calculate the timing differences between the individual trials and the corrected beat times.




```
trialdiffs = Transpose[Transpose[data] - Transpose[mcbeats][[2]]];







Map[Mean[Abs[QScore[#]]] &, trialdiffs]







{2.32884, 2.33993, 2.13148, 2.36285, 2.40196, 2.17492,
 2.15315, 2.37134, 2.2189, 2.22004, 2.74805, 2.48261, 2.10141,
 2.48097, 2.38752, 2.39842, 2.14267, 2.43396, 2.42107, 2.26832}
```

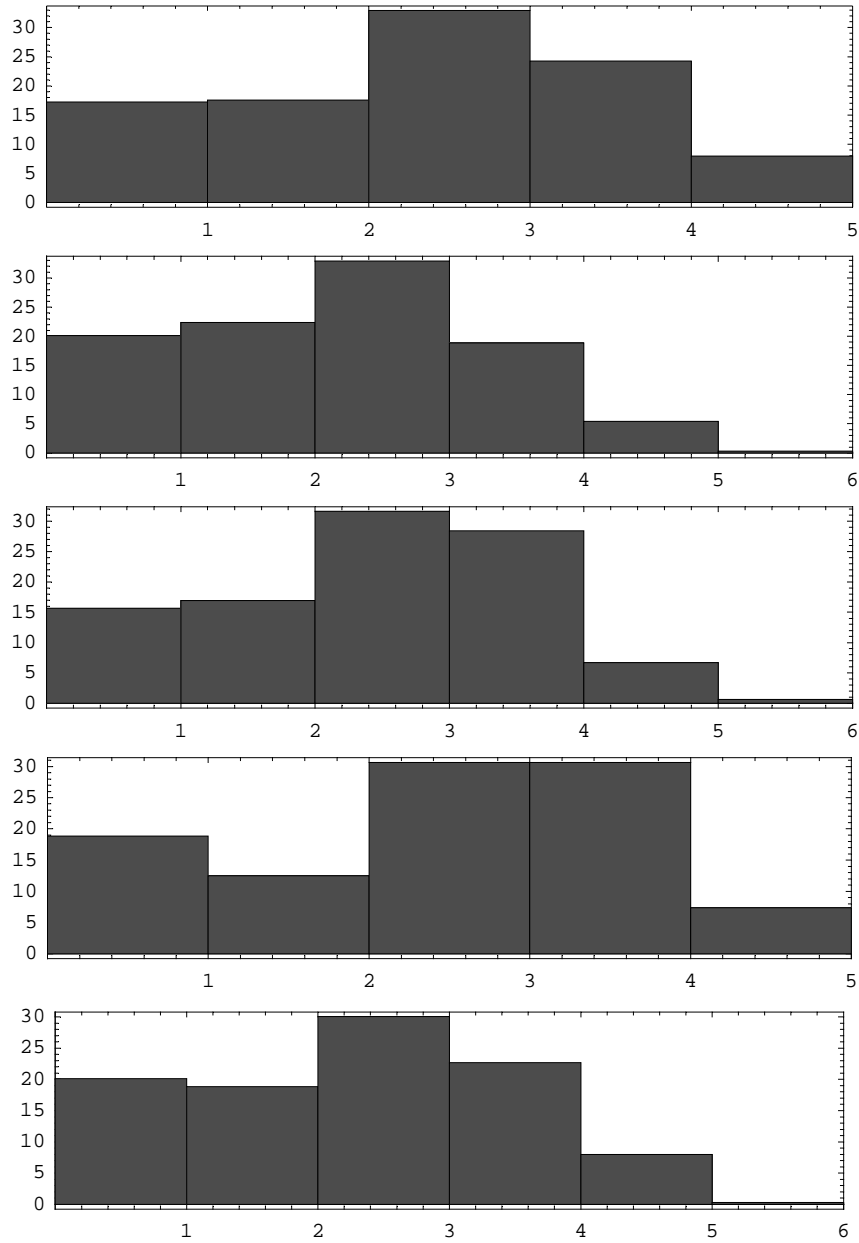
Here are quality score plots for each trial:

```
Table[Histogram[Abs[QScore[trialdiffs[[xz]]]], AspectRatio -> 1/4,
  Frame -> True, Range -> {{0, 6} {0, 30}}, HistogramScale -> 100,
  HistogramCategories -> 5], {xz, 1, Length[trialdiffs]}];
```










How does the accuracy of the first half of the trials compare to the second half of the trials?:


```
firsttrials = Drop[trialdiffs, -10];  
lasttrials = Drop[trialdiffs, 10];  
avgfirst = Map[Mean, Transpose[firsttrials]];  
avglast = Map[Mean, Transpose[lasttrials]];
```



```
Histogram[Abs[QScore[avgfirst]], AspectRatio → 1 / 4,
HistogramScale → 100, Frame → True, HistogramCategories → 5];
```



```
Histogram[Abs[QScore[avglast]], AspectRatio → 1 / 4,
HistogramScale → 100, Frame → True, HistogramCategories → 5];
```



```
Mean[Abs[QScore[avgfirst]]]
```

```
2.06826
```

```
Mean[Abs[QScore[avglast]]]
```

```
2.23328
```

On the average, beats in the second half of the trials are 6 milliseconds closer to the "true" beat locations:

```
firsterr = Mean[Abs[avgfirst]]
```

```
58.4745
```

```
lasterr = Mean[Abs[avglast]]
```

```
63.2942
```

```
firsterr - lasterr
```

```
-4.81971
```

```
allerr = N[Mean[Abs[mdiffs]]]
```

```
59.6454
```

```
allerr - lasterr
```

```
-3.64886
```


■ Plotting the average displacement error for each trial

This section plots the average displacement error for each trial which shows the gradual improvement in the reverse conducting over time.

```
avgerr = Table[Apply[Plus, Sqrt[ trialdiffs[[i]] * trialdiffs[[i]] ] ] /
  Length[trialdifs[[i]]] * 1.0 , {i, 1, Length[trialdifs]}]

{72.4633, 70.6297, 64.9777, 71.4505, 73.115, 63.3699, 64.138, 71.8934, 64.6752, 63.3419,
  88.0238, 76.6205, 61.7071, 75.7994, 71.4618, 71.0112, 60.5448, 73.2119, 72.1487, 67.128}


errplot = ListPlot[avgerr, PlotJoined → True, Frame → True, AspectRatio → 1 / 3];
```




Now fit an exponentially decreasing line though the trial errors to model what the "learning curve" is for reverse conducting this performance.

```
logfit = Fit[Log[10, avgerr], {1, x}, x]
1.84083 + 0.000179158 x

fitplot = Plot[10^logfit, {x, 0, 20}, PlotStyle → Hue[0]];
```



```
Show[errplot, fitplot, errplot];
```



Predictions of accuracy can be derived from the learning curve over time:

```
err40 = 10^logfit /. x -> 40
```

```
70.468
```

```
lasterr - err40
```

```
-7.17377
```

```
err100 = 10^logfit /. x -> 100
```

```
72.234
```

```
lasterr - err100
```

```
-8.93973
```

■ Effect of removing the "worst" trial

What happens if the "worst" trial is removed from the average tapping time for each beat?

```
maxindex = 1;
```

```
For[i = 1, i < Length[avgerr], i++,
```

```
  If[avgerr[[maxindex]] < avgerr[[i]], maxindex = i];
```

```
];
```

```
maxindex = maxindex
```

```
11
```

```
newerrset = RotateRight[Drop[RotateLeft[trialdiffs, maxindex], -1], maxindex - 1];
```

```
newavgerr = Table[
```

```
  Apply[Plus, Sqrt[newerrset[[i]] * newerrset[[i]]] / Length[newerrset[[i]]] * 1.0 ,
```


```
  {i, 1, Length[newerrset]}]
```

```
{72.4633, 70.6297, 64.9777, 71.4505, 73.115, 63.3699, 64.138, 71.8934, 64.6752, 63.3419,
```

```
76.6205, 61.7071, 75.7994, 71.4618, 71.0112, 60.5448, 73.2119, 72.1487, 67.128}
```

```
newplot =
```


```
ListPlot[newavgerr, PlotJoined -> True, Frame -> True, AspectRatio -> 1 / 3, Axes -> False];
```




```
logfit = Fit[Log[10, newavgerr], {1, x}, x]
```

```
1.83717 + 0.0000176669 x
```

```
newfitplot = Plot[10^logfit, {x, 0, 20}, PlotStyle -> Hue[0]];
```



```
Show[newplot, newfitplot, newplot];
```



```
newavg = Mean[Abs[Map[Mean, Transpose[newerrset]]]]
```

```
59.0263
```

```
allerr
```

```
59.6454
```

■ Trying to identifying "worst" performances from the average


Instead of comparing the individual performances to the corrected data, compare them to the average of all trials.

```
avgtrialdiffs = Transpose[Transpose[data] - Transpose[mabeats][[2]]];
```

```
errfromavg = Map[Mean, Abs[avgtrialdiffs]] // N
```

```
{33.3163, 31.4366, 34.6021, 29., 31.9936, 29.8034,  
34.6352, 31.7007, 34.7039, 31.2492, 42.7531, 30.9235, 40.1449,  
31.2338, 29.6853, 28.5919, 34.5137, 30.6515, 34.8707, 26.9371}
```

```
errminusone = ListPlot[errfromavg, PlotJoined -> True,  
Frame -> True, PlotRange -> All, Axes -> False, AspectRatio -> 1/3];
```



■ Dropping any single trial

Dropping the worst trial does not improve the accuracy of the average. Does dropping any other trial improve the average of all trials?


```
droppedsets = Table[RotateRight[Drop[RotateLeft[trialdiffs, i], -1], i], {i, 1, 20}];

droppedavg = Table[Mean[Abs[Map[Mean, Transpose[droppedsets[[i]] ]]]], {i, 1, 20}]

{59.714, 59.8883, 60.0786, 59.694, 59.7629, 60.0164,
 60.2486, 59.8035, 60.2439, 60.1125, 59.0263, 59.2182, 60.4516,
 59.5161, 59.6182, 59.5889, 60.4965, 59.4451, 59.5846, 59.7559}

plot = ListPlot[droppedavg, PlotJoined → True,
  Frame → True, Axes → False, DisplayFunction → Identity];

Show[plot, Graphics[{RGBColor[0, 1, 0], Line[{{0, allerr}, {20, allerr}}]}],
  DisplayFunction → $DisplayFunction, AspectRatio → 1 / 3];
```



■ Dropping a range of trials

Now examine how the displacement error changes as more and more of the earlier trials are removed from the average.

```
dsets = Table[Drop[trialdiffs, i], {i, 1, 19}];


davg = Table[Mean[Abs[Map[Mean, Transpose[dsets[[i]] ]]]], {i, 1, 19}]

{59.714, 59.8769, 60.2334, 60.2462, 60.304, 60.6605, 61.2898, 61.5655, 62.2953, 63.2942,
 62.2026, 61.2284, 63.0674, 62.7498, 62.2279, 61.7733, 66.2636, 66.6032, 67.128}

dplot = ListPlot[davg, PlotJoined → True,
  Frame → True, Axes → False, DisplayFunction → Identity];

dplotp = ListPlot[davg, PlotStyle → PointSize[0.02],
  Frame → True, Axes → False, DisplayFunction → Identity];
```

```
dropstart =
  Show[dplot, dplotp, Graphics[{RGBColor[0, 1, 0], Line[{{1, allerr}, {19, allerr}}]}],
  FrameTicks -> {Range[19], Automatic, None, None},
  DisplayFunction -> $DisplayFunction, PlotRange -> All, AspectRatio -> 1 / 3];
```



Here is a plot created by dropping the later trials more and more:


```
rdsets = Table[Drop[trialdiffs, -i], {i, 1, 19}];

rdavg = Table[Mean[Abs[Map[Mean, Transpose[rdsets[[i]]]]]], {i, 1, 19}]
{59.7559, 59.6568, 59.3437, 60.2326, 60.0429, 59.953, 59.6653, 60.8255, 60.169, 58.4745,
  58.9942, 60.1895, 60.1333, 61.8746, 63.7693, 64.5256, 64.098, 68.3389,  $\frac{22681}{313}$ }


rdplotp = ListPlot[rdavg, PlotStyle -> {RGBColor[0, 0, 1], PointSize[0.02]},
  Frame -> True, Axes -> False, DisplayFunction -> Identity];

rdplot = ListPlot[rdavg, PlotJoined -> True, Frame -> True,
  Axes -> False, DisplayFunction -> Identity, PlotStyle -> RGBColor[0, 0, 1]];

dropend = Show[rdplot, rdplotp,
  Graphics[{RGBColor[0, 1, 0], Line[{{1, allerr}, {19, allerr}}]}],
  FrameTicks -> {Range[19], Automatic, None, None},
  DisplayFunction -> $DisplayFunction, PlotRange -> All, AspectRatio -> 1 / 3];
```



```
Show[dropstart, dropend];
```




```
individ = Map[Mean[N[Abs[# - Transpose[mcbeats][[2]]]]] &, data]
```

```
{72.4633, 70.6297, 64.9777, 71.4505, 73.115, 63.3699, 64.138, 71.8934, 64.6752, 63.3419,
  88.0238, 76.6205, 61.7071, 75.7994, 71.4618, 71.0112, 60.5448, 73.2119, 72.1487, 67.128}
```


```
individplot = ListPlot[individ, PlotStyle → {Hue[0.0], PointSize[0.02]},
  Frame → True, Axes → False, DisplayFunction → Identity, AspectRatio → 1/3];
```

```
individplotp = ListPlot[individ, PlotJoined → True, Frame → True,
  Axes → False, PlotStyle → Hue[0.0], DisplayFunction → Identity];
```

```
Show[individplot, individplotp, DisplayFunction → $DisplayFunction];
```



```
Show[individplot, individplotp,
  dropstart, dropend, DisplayFunction → $DisplayFunction];
```



The red curve in the plot above shows the average displacement errors from the corrected times for each of the 20 individual trials. The blue curve represents dropping more and more of the later trials being dropped. The black curve represents dropping more and more of the earlier trials.

■ Offset sensitivity

Individual trials were aligned to the test case using a few points which seemed to be in a stable tempo region. How accurate is the offset calculated from these few points in the audio? Would the average of all trials improve if an offset is calculated for each trial based on the corrected data?

```

SquareDiff[x_, y_] := Apply[Plus, (x - y) * (x - y)]

AlignSequences[seq1_, seq2_] := Module[
  {diffs},
  diffs = Table[{x, SquareDiff[seq1, seq2 + x]}, {x, -10000, 10000, 100}];
  (x /. Solve[D[Fit[diffs, {1, x, x^2}, x], x] == 0, x])[1]
]

AlignSequences[{1, 2, 3}, {2, 3, 4}]

-1.

addoffsets =
  Table[AlignSequences[Transpose[mcbeats][[2]], data[[i]], {i, 1, Length[data]}]
  {-32.7572, -31.6668, -18.1894, -39.0607, -45.1342, -32.7307,
  -32.1319, -45.1974, -21.6061, -34.2045, -72.1559, -47.9067, -13.2294,
  -42.9977, -31.5312, -37.9211, -15.413, -46.0057, -45.6271, -30.945}

newmabeats = Map[Mean, Transpose[data + addoffsets]];

extraoffset = Mean[addoffsets]

-35.8206

Min[addoffsets]
Max[addoffsets]

-72.1559

-13.2294

omdiffs = Transpose[mcbeats][[2]] - newmabeats;

N[Mean[Abs[omdiffs]]]

50.7885

```

■ Comparing corrected beats to average reverse conducting beats

This section plots the average tempo range and compares it to the actual tempos measured from the audio file.

```

Get[performance <> ".ma", Path → "D:\mazurkas\eval"]
data = Transpose[ToExpression[pvar]];

```



```

eachbeatstats = Map[DispersionReport, data] // N;


confidence =
  Map[MeanCI[#[[1]], KnownStandardDeviation → (KnownStandardDeviation /. #[[2]])] &,
    Transpose[{data, eachbeatstats}]];

durationmean = Map[Mean, data] // N;
maxdurs = Map[Max, data];
mindurs = Map[Min, data];

meanhighplot = ListPlot[Transpose[confidence][[1]], PlotJoined → True,
  DisplayFunction → Identity, PlotStyle → RGBColor[.5, .5, .5]];
meanlowplot = ListPlot[Transpose[confidence][[2]], PlotJoined → True,
  DisplayFunction → Identity, PlotStyle → RGBColor[.5, .5, .5]];
maxplot = ListPlot[maxdurs, DisplayFunction → Identity,
  PlotJoined → True, PlotStyle → RGBColor[0.8, 0.8, 0.8]];
minplot = ListPlot[mindurs, DisplayFunction → Identity,
  PlotJoined → True, PlotStyle → RGBColor[0.8, 0.8, 0.8]];

meanplot = ListPlot[durationmean,
  PlotJoined → True, PlotRange → All, AspectRatio → 1 / 3, Frame → True,
  FrameLabel → {"beats", "milliseconds", "measures", "tempo"},
  FrameTicks → {Automatic, Automatic, Table[{x, x / 3}, {x, 0, 315, 3}],
    Map[{60000.0 / #, #] &, {30, 40, 44, 48, 54, 60, 70, 82, 100, 128, 180, 280}}]];


```



```

allplot =
  Show[meanplot, maxplot, minplot, meanhighplot, meanlowplot, meanplot, Axes → False];

```



```

correctdeltas = Drop[RotateLeft[corrected, 1] - corrected, -1];


```

```


cordels = Transpose[
  Partition[Transpose[{Drop[Transpose[mcbeats][[1]] + 1, -1], correctdeltas}], 3]];

```

```
correctplot = Show[Graphics[
  {PointSize[0.02], RGBColor[1, 0, 0], Map[Point, cordels[[1]]], RGBColor[0, 0, 1],
  Map[Point, cordels[[2]]], Map[Point, cordels[[3]]]}], AspectRatio -> 1 / 4];
```

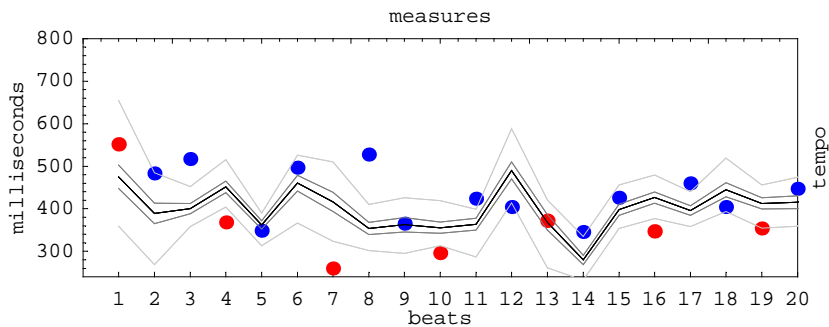


```
newallplot = Show[allplot, correctplot];
```



```
finalplot =
```

```
Show[allplot, correctplot, allplot, Frame -> True, PlotRange -> {{0, 20}, {240, 800}},
  FrameTicks -> {Range[20], Automatic, Automatic, None}];
```



```
ShowFrame[finalplot_, minrange_, maxrange_, len_, frame_, width_] := Module[
  {framecount, startpoint, endpoint, rmin, rmax, starti, endi, delta},
  framecount = len / width + 0.99 // Floor;
  startpoint = frame * width + 1;
  endpoint = (frame + 1) * width + 1;
  starti = (startpoint - 1) + 1;
  endi = (endpoint - 1) + 1;
  If[endi > len, endi = len];
  If[starti > len, starti = len];
  rmax = Max[Take[maxrange, {starti, endi}]];
  delta = rmax * 0.03;
  rmax = rmax + delta;
  rmax = 1500;
  rmin = Min[Take[minrange, {starti, endi}]];
  rmin = rmin - delta;
  rmin = 200;
  Show[finalplot, PlotRange -> {{startpoint, endpoint}, {rmin, rmax}},
  FrameLabel -> {"beats", "milliseconds", "measures", "tempo"},
  FrameTicks -> {Automatic, Automatic, Table[{x + 1, x / 3 + 1}, {x, 0, 2000, 3}],
  Map[{60000.0 / #, #} &, {30, 40, 44, 48, 54, 60, 70, 82, 100, 128, 180, 280}]},
  Axes -> False]
]
```

Here is a zoom in on the tempo for every 8 measures. The black line indicates the mean reverse-conducted durations for every beat. The dark gray lines surrounding the average duration line is the 95% confidence range for the true mean, and the light gray lines indicate the maximum and minimum durations for each beat from all reverse conducting trials.

The red and blue dots indicates the beat event durations from the audio file which are more accurate than the average reverse conducting durations, and can be assumed to be the "correct" answer. The red dots represent the first beat of a measure, and the blue dots represent the other beats in the measure.

```
Table[ShowFrame[newallplot, mindurs, maxdurs, Length[data], i, 24],
  {i, 0, Floor[Length[data] / 8.0 / 3.0 + .99] - 1}];
```

