

How (not) to assess the importance of correlations for the matching of spontaneous and evoked activity

József Fiser¹, Máté Lengyel², Cristina Savin², Gergő Orbán^{2,3}, Pietro Berkes⁴

¹Department of Cognitive Science, Central European University, ²Computational and Biological Learning Lab, Department of Engineering, University of Cambridge, ³Wigner Research Institute for Physics, Hungarian Academy of Sciences, ⁴Volen National Center for Complex Systems, Brandeis University

Contents

February 2013:	Comment on Okun et al. (2012)	p.2
March 2013:	Response to the reply of Okun et al. (2013) to our comments	p.6

Comment on Okun et al. (2012)

In a recent publication, Okun et al. (2012) criticise our earlier paper (Berkes et al., 2011) in which we analysed multi-neuron firing patterns in V1 of awake ferrets. We described a remarkably close match between the distribution of multi-neuron firing patterns during spontaneous activity (recorded in the absence of visual stimulation) and during evoked activity recorded while stimulating with naturalistic visual stimuli. We argued that this match confirmed predictions about the cortex implementing a statistically optimal internal model of the environment. In contrast, Okun et al. claim that this match can be a consequence of trivial statistical properties of multi-neuron firing patterns, specifically fluctuations in overall population firing rates, and as such these epiphenomenal results are not indicative of optimal internal models. Below we explain that the analyses conducted by Okun et al. suffer from both conceptual and statistical flaws which invalidate their interpretation of their own data. We also show that if the correct analyses are performed on our original data set, the claims of Okun et al. regarding our findings do not hold and in fact these analyses provide additional support for our main conclusion.

Following the logic of Berkes et al., central to Okun et al.'s analysis is the construction of a surrogate data set that respects as constraints some simple statistical properties of the experimentally measured distribution of multi-neuron firing patterns, but is otherwise as random as possible (it has maximum entropy given the constraints). One can then test whether there is a difference between the original results obtained with the real data and the results obtained using surrogate data. If no significant difference can be found, then the original results are deemed epiphenomenal. Berkes et al. used the simplest possible such surrogate: a distribution under which all neurons fire independently from each other at their true observed firing rates – this has been a standard test for the role of correlations (of 2nd and higher order) in shaping the distribution of neural responses since the seminal paper of Schneidman et al. (2006). Berkes et al. reported that this surrogate data set did not capture some essential aspects of the real neural responses (Berkes et al., Fig. 3A-B) and thus that their results could not be trivially explained by single neuron firing rate dynamics.

Building on recent advances in the field of maximum entropy analyses of neural data (Tkačik et al., 2012), Okun et al. go one step further and construct a more sophisticated surrogate, which preserves not only single neuron firing rates, but also the distribution of the number of co-active units at a time. This is an important control because, in general, this second constraint introduces correlations at all orders, without any functionally relevant pairwise coupling between neurons. Indeed, the number of degrees of freedom of this surrogate (a measure of its complexity) is still only linear in the number of units, as it is in the simple surrogate that only cares about single neuron firing rates, while in a true second-order surrogate (Schneidman et al., 2006), that is able to capture the detailed pairwise (functional) connections between neurons, this measure scales quadratically with the number of units. The surrogate used by Okun et al. is also appealing because it allows a simple biological interpretation: a network in which units are functionally disconnected (or randomly connected) but undergo synchronized fluctuations of activity would produce data whose statistics would match those of this surrogate (see Fig. 8 in Okun et al.).

Okun et al. then perform the following key analyses using this surrogate data:

1. They show that the surrogate data is very close to the true data recorded in A1 (Okun et al., Fig. 3).
2. They replace both the true spontaneous and evoked activity distributions with their surrogate counterparts and show that some of the results of Berkes et al. (Berkes et al., Figs. 2B, 3A-B) can be replicated: the statistical dissimilarity between these surrogates can be minimal in adult animals both in A1 and V1 (Okun et al., Figs. 4B and 5) and substantially decrease over (simulated) development (Okun et al., Figs. 7 and 8).

These analyses have questionable relevance for the earlier results of Berkes et al. on both counts:

1. Assessing match between surrogate and true data.
 - a) Okun et al. never demonstrate directly a match between surrogate and true data for their V1 recordings. Given that all the rest of their analyses pertaining directly to the results of Berkes et al. (obtained in V1) are performed on the V1 data set (Okun et al., Figs. 7-8), demonstrating this match would have been a necessary prerequisite for validating their approach.

- b) In fact, even for their A1 data, Okun et al.’s own analyses do indicate a divergence between the surrogate and true data that appears to be significant (Okun et al., Fig. 3B, difference between gray and red bars). This significance is never tested statistically, even though if there was a significant divergence that would invalidate their claims about their surrogate reproducing the relevant statistics of the true data (as is the case in their Fig. 4C, see below in point 2b).
 - c) Repeating their analysis with their surrogate on the data set used in Berkes et al., we find that the divergence between real and surrogate data significantly increases over development (Spearman’s $\rho = 0.68, p = 0.005$ for spontaneous, and Spearman’s $\rho = 0.57, p = 0.027$ for movie-evoked activities) and reaches highly significant levels in adult age groups (Fig. 1A).
2. Using surrogate-to-surrogate divergences.
- a) Measuring the divergence between surrogate spontaneous and surrogate evoked activities, as Okun et al. did, is prone to trivial outcomes. Imagine that the original spontaneous and evoked activities are truly being matched by all statistical measures. Then it follows, by definition, that their corresponding surrogate versions, which respect only a subset of these statistical measures, will also be automatically matched to each other.* (This is strictly true only if the divergence between the true distributions is zero, but will usually hold for small but non-zero divergences as well.) As expected, this effect can be demonstrated in the data set of Berkes et al. using Okun et al.’s surrogate: the divergence between surrogate spontaneous and surrogate movie-evoked activity significantly decreases over development (Spearman’s $\rho = -0.75, p = 0.0014$) and reaches non-significant levels in adult animals (Fig. 1B, purple bars). To illustrate that this effect is indeed trivial, we computed the same divergence with the simple surrogate used by Berkes et al. respecting only single-neuron firing rates, and found qualitatively identical results (Spearman’s $\rho = -0.71, p = 0.003$, Fig. 1B, pink bars).
 - b) The appropriate way to establish the importance of higher-order statistics for matching two histograms is to perform analyses of the kind shown in Figure 3B of Berkes et al.: surrogate spontaneous must be compared to true evoked activities (or vice versa). If it is only the lower order moments reproduced by the surrogate that matter for the original match between the true distributions, then this divergence using the surrogate should be similar to that computed between the true distributions. By inversion, a substantially larger surrogate-to-true divergence than true-to-true divergence would indicate that higher order correlations were important for the original true-to-true divergence. Okun et al. show only a single example of such an analysis using their A1 (but not V1) data (Okun et al., Fig. 4C) and even this analysis is inconclusive at best. According to their conclusion from this analysis, not supported by any statistical test, there is no substantial difference between surrogate-to-true and true-to-true divergences – even though their figure suggests that there is (the green dots are visibly above the diagonal). In fact, Okun et al. themselves acknowledge that the difference they found in their divergences is due to the inappropriate fit of their surrogate to their real data in the first place (Okun et al., Fig. 3B, see point 1b above). This assessment seems to be correct, but it also points to a crucial flaw undermining their ensuing claims as explained above: when surrogate-to-true divergences are larger than true-to-true divergences then higher-order correlations do play a major role.
 - c) Most importantly, when the correct surrogate-to-true analysis is performed on our data, we find that in contrast to true-to-true divergences (Fig. 1C, red bars), surrogate-to-true divergences do not decrease significantly over time (Spearman’s $\rho = -0.32, p = 0.246$) and remain significantly above baseline even in adult animals (Fig. 1C, purple bars), further strengthening Berkes et al.’s earlier results obtained using the simpler surrogate analysis (Fig. 1C, pink bars).

In summary, we concur with the spirit of Okun et al.’s work: showing a low (and developmentally decreasing) divergence between spontaneous and evoked activities is necessary but not sufficient for claiming statistical

*As an example, imagine that there are two data sets, each with complicated non-Gaussian looking histograms. We create a Gaussian surrogate for each, matching their mean and variance. If the original histograms looked the same in all their gory detail, their Gaussian surrogates will also look the same – but this does not mean that the original histograms didn’t match in ways that went beyond their mean and variance (Fig. 2A).

optimality. In fact, this is in line with the approach already taken by Berkes et al. who provided a set of controls specifically to address this issue, including controls for the effects of spatial correlations, temporal correlations, and stimulus statistics. Okun et al. elaborate on one of these controls, concerning spatial correlations. Indeed, the surrogate they proposed as an improved control does allow for a more stringent test of trivial factors affecting spatial correlations than the simple surrogate normally used (Schneidman et al., 2006; Berkes et al., 2011). In that respect, their paper represents an important contribution which will be a useful addition to the toolbox of maximum entropy models in systems neuroscience (Schneidman et al., 2006; Tkačik et al., 2012).

Unfortunately, however, there are two reasons why the conclusions Okun et al. drew from their otherwise useful surrogate are questionable. First, the way they ended up using this surrogate in their analyses was mostly inappropriate to the main question they tried to address, and thus provided trivial results. (And in those cases in which their analysis was appropriate, in their Fig. 4C, their results seem to agree with those reported in Berkes et al.) Second, we performed the correct analyses with their surrogate on our data, and found that these provide additional support for the main conclusion of Berkes et al.: the match between spontaneous and movie-evoked activities in the visual cortex is remarkable, cannot be explained by trivial statistical properties and it is, therefore, indicative of statistically optimal internal models in the cortex.

References

- Berkes P, Orbán G, Lengyel M, Fiser J (2011) Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* 331:83–87.
- Okun M, Yger P, Marguet SL, Gerard-Mercier F, Benucci A, Katzner S, Busse L, Carandini M, Harris KD (2012) Population rate dynamics and multineuron firing patterns in sensory cortex. *J Neurosci* 32:17108–19.
- Schneidman E, Berry, 2nd MJ, Segev R, Bialek W (2006) Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440:1007–12.
- Tkačik G, Marre O, Mora T, Amodei D, Berry, 2nd MJ, Bialek W (2012) The simplest maximum entropy model for collective behavior in a neural network. *arXiv preprint arXiv:1207.6319*.

An abbreviated version of this text has also been posted as a Letter to the Editor of the *Journal of Neuroscience* as a reply to Okun et al. (2012).

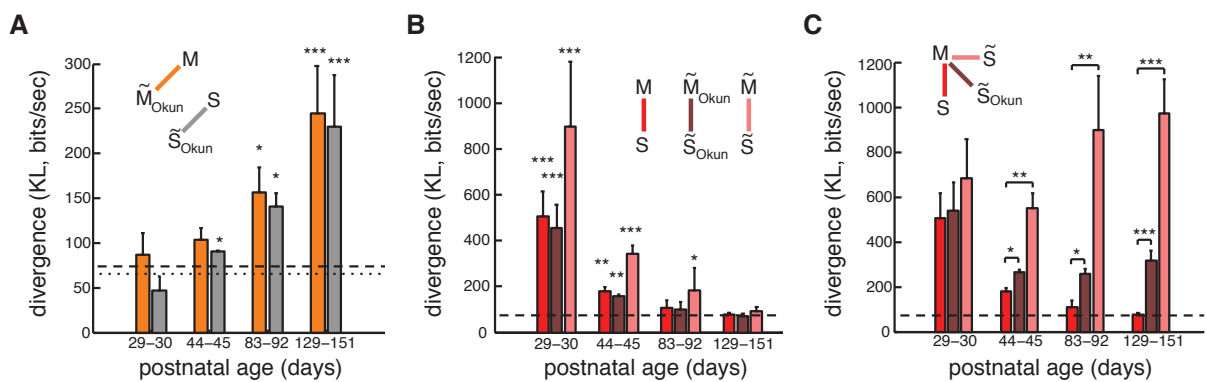


Figure 1.

A. Divergences between true activity distributions and their surrogate counterparts. As a reference, the dashed and dotted lines show the average of the within-condition baselines (dotted for S, dashed for M), computed with within-condition data split into two halves.

B. Divergences between movie-evoked (M) and spontaneous (S) activities using the true data (red bars), the surrogate used by Okun et al. respecting single neuron firing rates and population rate distributions (purple bars), and the original surrogate used by Berkes et al. respecting only single neuron firing rates (pink bars). Red bars and dashed line are reproduced from Berkes et al.

C. Divergences between true movie evoked and the true (red bars) surrogate spontaneous activities (purple bars: Okun et al.'s surrogate, pink bars: Berkes et al.'s surrogate). Red and pink bars and dashed line (computed as in panel A) are reproduced from Berkes et al.

Asterisks denote significant differences from respective baselines (A, B) or between bars (C), * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, m-test (for details see Berkes et al.).

Response to the reply of Okun et al. (2013) to our comments

Following our comment (above) on their original paper (Okun et al., 2012), Okun et al. (2013) recently published a reply. Our summary for where this discussion stands is the following. There are two general points that Okun et al. (2012, 2013) made about how the learning of internal models could and should be tested in experiments, which we agree with in principle, although with some important qualifications. There are also two important specific points about their own results and analyses of data, and the way they interpreted those results, that we disagree with and still find technically flawed. In our view, these flaws will need to be corrected and some of the main conclusions of Okun et al. (2012) will need to be revisited in light of these analyses. We detail these points below.

Agreement on general points

Looking at KL divergences (or any other similar measure of divergence) between spontaneous and evoked activities alone cannot be taken as conclusive evidence that an internal model of the environment has been brought about by learning. We could not agree more and have never in fact said otherwise. In our original paper, we wrote that our findings “do not address the degree to which statistical adaptation in the cortex is driven by visual experience or by developmental programs” (Berkes et al., 2011), and thus interpreted our results as a sign of adaptation (learning-driven or otherwise), not learning *per se*.[†] Thus, Okun et al.’s statement that our result alone “does not indicate that an environmental model has been learned” echoes our earlier views. However, we now also have results from newer experiments using lid-sutured animals that had no normal vision until the moment of the tests (Fiser et al., 2012; Savin et al., 2013). These results, to which Okun et al. also refer and which we will present in detail elsewhere, have since provided additional evidence clearly indicating, at least in the oldest age group, that visual experience does have a significant role in the specificity of the match between evoked and spontaneous activities. This is compatible with stimulus specificity only emerging in this age group in control animals (Berkes et al. (2011), Fig. 4). Strangely, Okun et al. interpret these results of ours as evidence for specifically non-learning related changes driving adaptation in all but the last age group. However, it is important to note that even in the first three age groups our results do not exclude forms of learning that our manipulation (lid suture) did not affect (such as cross-modal learning, learning from low spatial frequencies, etc), or our data recording and analyses techniques were unable to pick up. Thus, Okun et al.’s interpretation is not quite warranted.

It is important to control for the effects of populations rate fluctuations in maximum entropy model-based analyses of the role of correlations. We agree with this statement, too. To make the historical context clear: at the time of the publication of Berkes et al. (2011), the “standard” way of demonstrating the importance of correlations in maximum entropy analyses was to follow Schneidman et al. (2006) and compare the true data to the simple surrogate that only controls for single neuron firing rates but not for the population rate distribution. This is precisely the practice we followed in Berkes et al. (2011). The contribution of Okun et al. (2012) and others (Tkačik et al., 2012) is that they suggested a more stringent test, controlling for the effects of population rate fluctuations, which warrants the revisiting of *all* earlier results that used just the simpler surrogate. We have now revisited our results published in Berkes et al. (2011), and confirmed them (see above, Fig. 1), and it would be interesting to see which other earlier results (such as Schneidman et al., 2006, and, importantly, the results of Okun et al., 2012 themselves, as we argue below) survive this new, more stringent test.

Disagreement on specific results and interpretations in Okun et al. (2012)

Population rate fluctuations account for most of the seemingly correlation-driven effects in the data of Okun et al. (2012) which therefore implies that they offer an alternative explanation of the results of Berkes et al. (2011). In our comment on their paper (see above) we stated that this conclusion by Okun et al. (2012) rests on performing the wrong comparison of divergences that were computed with their otherwise valuable new surrogate. In their reply, Okun et al. (2013) dismiss our criticism by stating that “the point [of their paper, Okun et al.,

[†]A simple content analysis of Berkes et al. (2011) shows 13 occurrences of words with a root “adapt” and 0 occurrence of words starting with “learn”.

2012] was not that P and \hat{P}^\ddagger are identical” and that “the relationship of $D[P||\hat{Q}]$ and $D[P||Q]$ was not central [to their analysis]”. However, as we said before and explain in more detail below in a separate section, this is precisely the shortcoming of their approach that should be mended: analyzing those relationships would have been crucial to support their claims about the (un)importance of true pairwise correlations, and the analysis they chose instead to perform (comparing $D[P||Q]$ with $D[\hat{P}_{Okun}||\hat{Q}_{Okun}]$) is not only much more indirect in general than the analyses we suggested and performed on our data (computing $D[P||\hat{P}_{Okun}]$ and also comparing $D[P||Q]$ with $D[P||\hat{Q}_{Okun}]$), but in this case it also leads to wrong conclusions. Importantly, while we have taken seriously the issues they had raised and re-analyzed our data using the surrogate they suggested, thus confirming the crucial (and developmentally increasing) importance of true pairwise (and higher order) correlations for the matching of spontaneous and evoked activities in our data set (see above), Okun et al. did not perform the analyses we pointed out were necessary on their data set. As a result, judging from the analyses they have conducted so far, there is no indication that the important effects of true pairwise correlations shown in our data set are not present in theirs – although the final word can only be said once they perform the correct analyses on their data set.

A random network, such as that shown in Okun et al. (2012), could reproduce the key results of Berkes et al. (2011) When making this claim, Okun et al. continue to ignore one of the most important controls we already presented in Berkes et al. (2011). In fact, precisely to address this issue, we performed in our original paper (Berkes et al., 2011) additional tests using evoked activities obtained using different kinds of artificial stimulus ensembles (block noise or drifting gratings), and showed that the divergence of these evoked activities from spontaneous activity was significantly larger than the original divergence obtained using natural image movie-evoked activities (see Figure 4 therein). This shows that the cortex is indeed specifically adapted for natural stimuli. Given the importance of this control, it is puzzling why nowhere in Okun et al.’s original paper (2012) or in their latest reply (2013) there is any reflection on how their random network fares on it. Therefore, in lack of evidence otherwise, we are led to believe that this control clearly rules out the random-network account of our results proposed by Okun et al.

As a final conclusion, before we present our detailed argument about the right comparison of divergences below, we want to stress again that in our view the challenge is now for Okun et al. (2012) to correct these technical flaws by using the right divergences and comparisons on the data they recorded, and by performing all necessary tests on their simulations, before they draw conclusions about their own results and those presented in Berkes et al. (2011).

Detailed comments on the right comparison of divergences

In their reply to our comment, Okun et al. (2013) argue that instead of performing the analyses we suggested and performed on our data (computing $D[P||\hat{P}_{Okun}]$ and also comparing $D[P||Q]$ with $D[P||\hat{Q}_{Okun}]$), the comparison of $D[P||Q]$ with $D[\hat{P}_{Okun}||\hat{Q}_{Okun}]$ is the relevant analysis to perform.[§] In their eyes, the fact that the difference between these divergences is small supports their claim that “word distributions primarily reflect changes in population rate dynamic”. This is despite the fact that they agree with us that a finding that both divergences are zero (or close to it) is uninformative as to the importance of correlations (or in general, statistics not controlled by maximum entropy surrogates, Fig. 2A). In contrast, note that the divergence we suggest should be used, $D[P||\hat{Q}_{Okun}]$, is able to dissect the role of higher-order statistical structure in causing a match between two distributions in this case (Fig. 2A).

Interestingly, Okun et al. (2013) insist that when the individual divergences, $D[P||Q]$ and $D[\hat{P}_{Okun}||\hat{Q}_{Okun}]$, are significantly larger than zero, comparing them is still meaningful. Unfortunately, this is simply not the case: in general, comparing $D[P||Q]$ with $D[\hat{P}||\hat{Q}]$ (where \hat{X} is some maximum entropy surrogate of X) can be very misleading for two reasons. First, because the way differences between different kinds of statistics (e.g. moments) of two distributions jointly determine the total Kullback-Leibler (KL) divergence between two distributions is highly

[‡]Where their \hat{P} is \hat{P}_{Okun} in our notation, and P and Q can stand for movie-evoked (M) or spontaneous (S) activities.

[§]To be precise: for their A1 data they did actually perform the tests we propose, but still drew conclusions from the other test, which we think is misleading, for their V1 data they did not even perform the tests that would have been necessary.

non-linear. Second, because the KL divergence is not a proper metric. Fig. 2B shows two examples in which both $D[P\|Q]$ and $D[\tilde{P}\|\tilde{Q}]$ are large and are very close to each other. However, in one case (Fig. 2B, top) they only differ in their lower order moments (which are being controlled by the maximum entropy surrogates), while in the other case (Fig. 2B, bottom) they differ both in these lower order moments and in higher order moments (that are not controlled by the maximum entropy surrogates). In general, it can also be seen (Fig. 2C) that $D[P\|Q]$ varies non-monotonically as the differences between P and Q in higher-order statistics change, such that any arrangement of $D[P\|Q]$ being greater or smaller than $D[\tilde{P}\|\tilde{Q}]$ is possible. Therefore, it is not at all necessarily true that the difference of these two divergences are zero when their higher-order statistics match. This demonstrates that the fact that $D[P\|Q]$ and $D[\tilde{P}\|\tilde{Q}]$ are large and nearly identical, may be a reflection of “a property specific to word distributions recorded in the sensory cortex” (Okun et al., 2013), but it tells us nothing about the role of higher-order statistical structure in causing a mismatch between P and Q – which is what the original claim of Okun et al. (2012) was about. Note that our suggested divergence cannot resolve this issue either (Fig. 2B), and consequently, we never made any conjectures from the large-divergence case. As a side note, it would be interesting to see whether there are other divergence-comparisons that may be useful in this case.

In summary, while the comparison Okun et al. advocate is unable to resolve, in general, the role of correlations in the case of either a match or a mismatch between evoked and spontaneous activities, the comparison we proposed and used throughout can resolve this role in a match (but not in a mismatch). Thus, while using the correct comparison we have convincingly demonstrated that in our data set, correlations do play a major role in shaping neural activities and causing a match between evoked and spontaneous activities, Okun et al. still need to perform these analyses, especially on their V1 data set. Until these analyses are performed, their claims about the (un)importance of correlations in their data set, including the schematic of the relative distances between distributions they presented (Fig.1 in Okun et al., 2013) remains unjustified and potentially misleading. In fact, looking at the correct divergences in our data set, the version of their schematic would look rather different in adult animals: P and Q would be close ($D[M\|S] = 100$, Fig. 1B and C red), so would be \tilde{P}_{Okun} and \tilde{Q} ($D[\tilde{M}_{\text{Okun}}\|\tilde{S}_{\text{Okun}}] = 100$, Fig. 1B purple), but the cross-distances would be much larger ($D[M\|\tilde{M}_{\text{Okun}}] = 200$, Fig. 1A orange; $D[S\|\tilde{S}_{\text{Okun}}] = 200$, Fig. 1C purple; $D[M\|\tilde{S}_{\text{Okun}}] = 300$, Fig. 1C purple).

References

- Berkes P, Orbán G, Lengyel M, Fiser J (2011) Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* 331:83–87.
- Fiser J, Savin C, Berkes P, Chiu C, Lengyel M (2012) Visual experience drives the increase in similarity between spontaneous and stimulus evoked activity in V1. In: Society for Neuroscience Meeting.
- Okun M, Yger P, Harris KD (2013) How (not) to assess the importance of correlations for the matching of spontaneous and evoked activity: a response. *arXiv preprints arXiv:1303.3287*.
- Okun M, Yger P, Marguet SL, Gerard-Mercier F, Benucci A, Katzner S, Busse L, Carandini M, Harris KD (2012) Population rate dynamics and multineuron firing patterns in sensory cortex. *J Neurosci* 32:17108–19.
- Savin C, Berkes P, Fiser J, Lengyel M (2013) Similarity between spontaneous and sensory-evoked activity does suggest learning in the cortex. In: Cosyne.
- Schneidman E, Berry, 2nd MJ, Segev R, Bialek W (2006) Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440:1007–12.
- Tkačik G, Marre O, Mora T, Amodei D, Berry, 2nd MJ, Bialek W (2012) The simplest maximum entropy model for collective behavior in a neural network. *arXiv preprint arXiv:1207.6319*.

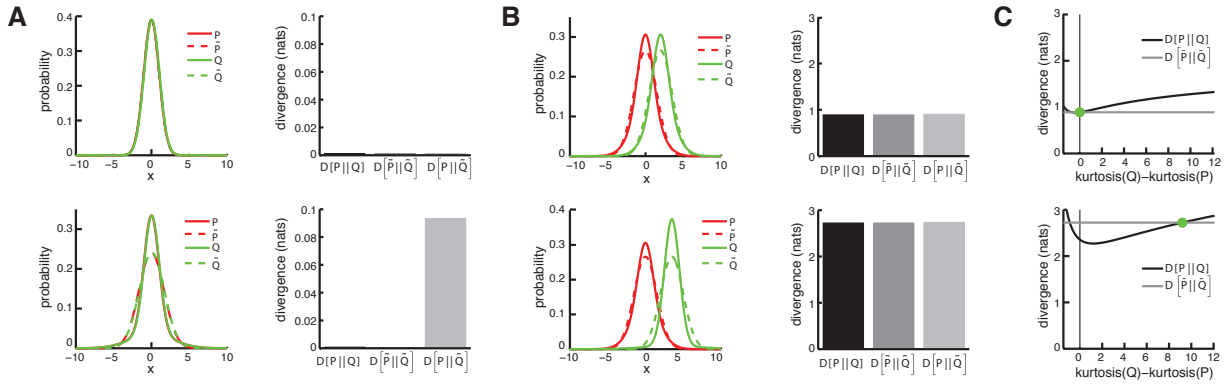


Figure 2.

A. When both $D[P||Q]$ (right panels, black) and $D[\tilde{P}||\tilde{Q}]$ (right panels, dark grey) are small and they are nearly identical, this does not distinguish between whether the match between P and Q is only governed by their lower order statistics (top) or also by their higher order statistics (bottom). Note that using $D[P||\tilde{Q}]$ (right panels, light grey) it is possible to distinguish between these two cases. In the examples shown, P and Q (left panels, red and green solid lines) are Gaussian scale mixtures (with two components), and their surrogates (\tilde{P} and \tilde{Q} , left panels, red and green dotted lines) are moment-matched normal distributions. Panels on the right show respective divergences. True distributions for the top panels are almost exactly Gaussians, hence they are near-identical to their surrogates. True distributions for the bottom panels are leptokurtotic and hence are not identical to their surrogates. P and Q are identical in both panels except for a very small shift in their means.

B. When both $D[P||Q]$ and $D[\tilde{P}||\tilde{Q}]$ are large and they are nearly identical, this still does not distinguish between whether the mismatch between P and Q is only governed by their lower order statistics (top) or also by their higher order statistics (bottom). Note that $D[P||\tilde{Q}]$ also does not resolve this issue. Lines in left panels and bars in right panels as in A. For the top panels, P and Q differ only in their means. For the bottom panels, P and Q also differ in their kurtosis (but not their variance or skewness) which is not controlled by their surrogates.

C. Lines show $D[P||Q]$ (black) and $D[\tilde{P}||\tilde{Q}]$ (dark grey) as the kurtosis of Q is varied relative to that of P , while keeping its mean, variance, and skewness fixed. Bottom and top panels as in B, green points show the kurtoses of the Q distributions used there. Note that in the case shown in the bottom panel a substantial difference in kurtosis between P and Q is accompanied by zero difference in the two divergences.