Fluorescence Imaging Analysis: The Case of Calcium Transients.

Christophe Pouzat

June 26 2013

Outline

The variability inherent to fluorescence imaging data (1)

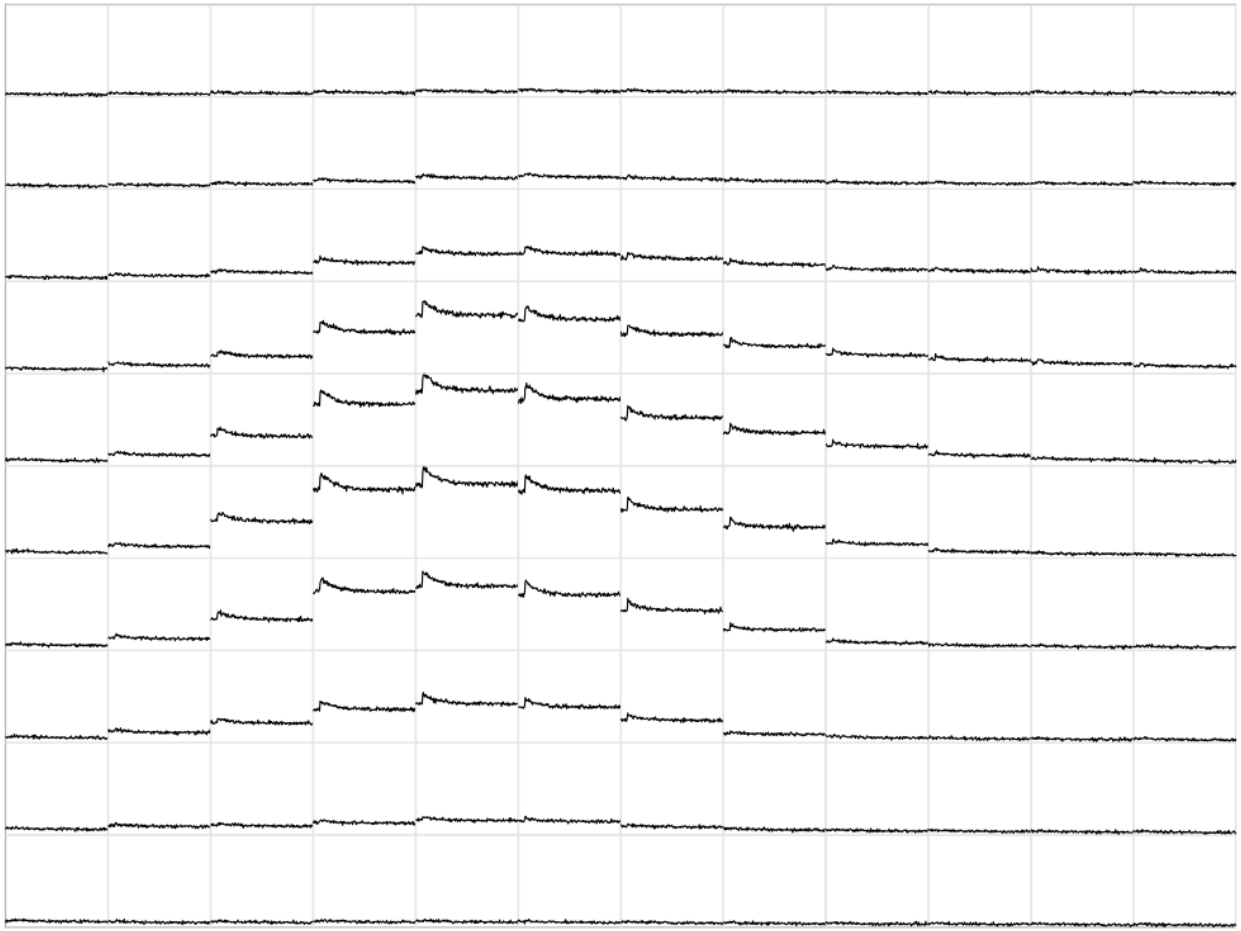

ADU counts (raw data) from Fura-2 excited at 340 nm. Each square corresponds to a pixel. 25.05 s of data are shown. Same scale on each sub-plot. Data recorded by Andreas Pippow (Kloppenburg Lab. Cologne University).

The variability inherent to fluorescence imaging data (2)

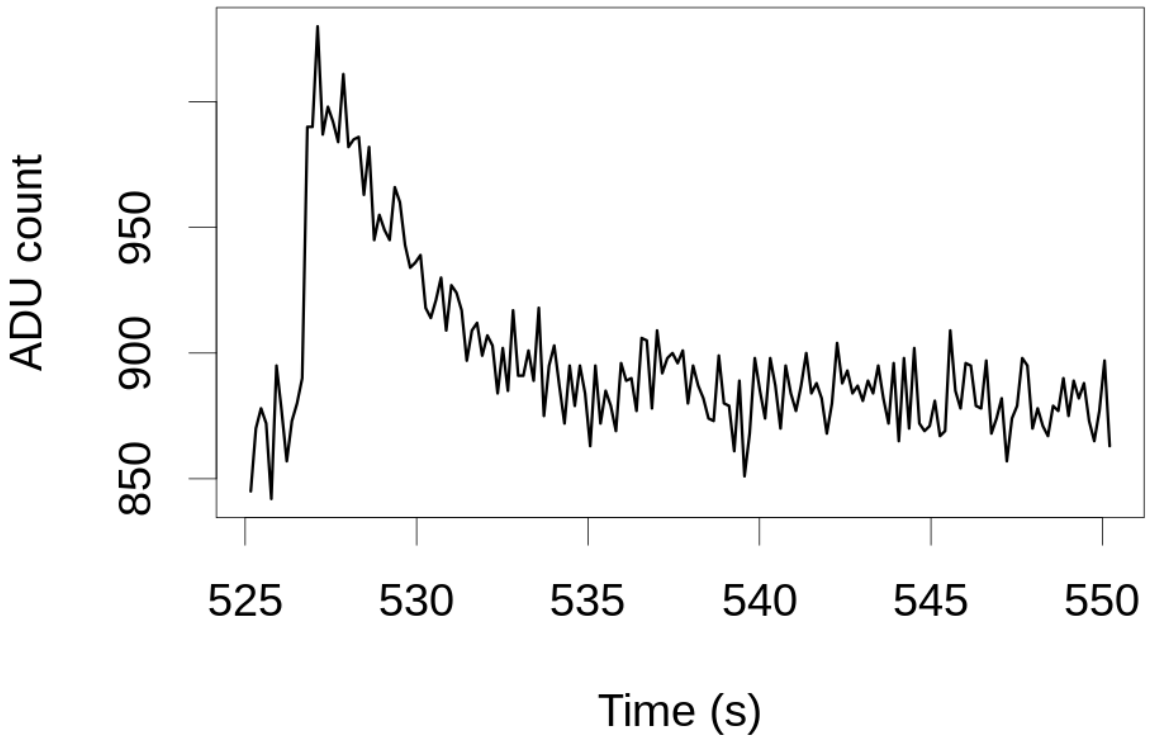

One of the central pixels of the previous figure.

What do we want? (1)

Given the data set illustrated on the last two slides we might want to estimate parameters like:

- the peak amplitude
- the decay time constant(s)
- the baseline level
- the whole time course (strictly speaking, a function).

What do we want? (2)

If we have a model linking the calcium dynamics—the time course of the free calcium concentration in the cell—to the fluorescence intensity like:

$$\frac{\mathrm{d}Ca_t}{\mathrm{dt}}\left(1+\kappa_F(Ca_t)+\kappa_E(Ca_t)\right)+\frac{j(Ca_t)}{v}=0\,,$$

where Ca_t stands for $[Ca^{2+}]_{free}$ at time t, v is the volume of the neurite—within which diffusion effects can be neglected—and

$$j(Ca_t)\equiv\gamma(Ca_t-Ca_{steady})\,,$$

is the model of calcium extrusion—Ca_{steady} is the steady state $[Ca^{2+}]_{free}$—

$$\kappa_F(Ca_t)\equiv\frac{F_{total}\,K_F}{(K_F+Ca_t)^2}\quad\text{and}\quad\kappa_E(Ca_t)\equiv\frac{E_{total}\,K_E}{(K_E+Ca_t)^2}\,,$$

where F stands for the fluorophore en E for the *endogenous* buffer.

What do we want? (3)

In the previous slide, assuming that the fluorophore (Fura) parameters: F_{total} and K_F have been calibrated, we might want to estimate:

- the extrusion parameter: γ
- the endogenous buffer parameters: E_{total} and K_E

using an equation relating measured fluorescence to calcium:

$$Ca_t = K_F \frac{S_t - S_{min}}{S_{max} - S_t},$$

where S_t is the fluorescence (signal) measured at time t, S_{min} and S_{max} are *calibrated* parameters corresponding respectively to the fluorescence in the absence of calcium and with saturating $[Ca^{2+}]$ (for the fluorophore).

What do we want? (4)

- The variability of our signal—meaning that under replication of our measurements *under the exact same conditions* we wont get the exact same signal—implies that our estimated parameters will also fluctuate upon replication.
- Formally our parameters are modeled as *random variables* and it is not enough to summarize a random variable by a single number.
- If we cannot get the full distribution function for our parameters, we want to give at least ranges within which the true value of the parameter should be found with a given probability.
- In other words: an analysis without confidence intervals is not an analysis, it is strictly speaking useless since it can't be reproduced—if I say that my time constant is 25.76 ms the probability that upon replication I get again 25.76 is essentially 0; if I say that the actual time constant has a 0.95 probability to be in the interval [24,26.5], I can make a comparison with replications.

A proper handling of the "variability" matters (1)

Let us consider a simple data generation model:

$$Y_i \sim \mathcal{P}(f_i)\,, \quad i = 0, 1, \ldots, K\ ,$$

where $\mathcal{P}(f_i)$ stands for the *Poisson distribution* with parameter f_i :

$$\Pr\{Y_i = n\} = \frac{(f_i)^n}{n!} \exp(-f_i)\,, \quad \text{for} \quad n = 0, 1, 2, \ldots$$

and

$$f_i = f(\delta i | f_\infty, \Delta, \beta) = f_\infty + \Delta\ \exp(-\beta\,\delta i)\ ,$$

δ is a time step and f_∞, Δ and β are model parameters.

A proper handling of the "variability" matters (2)

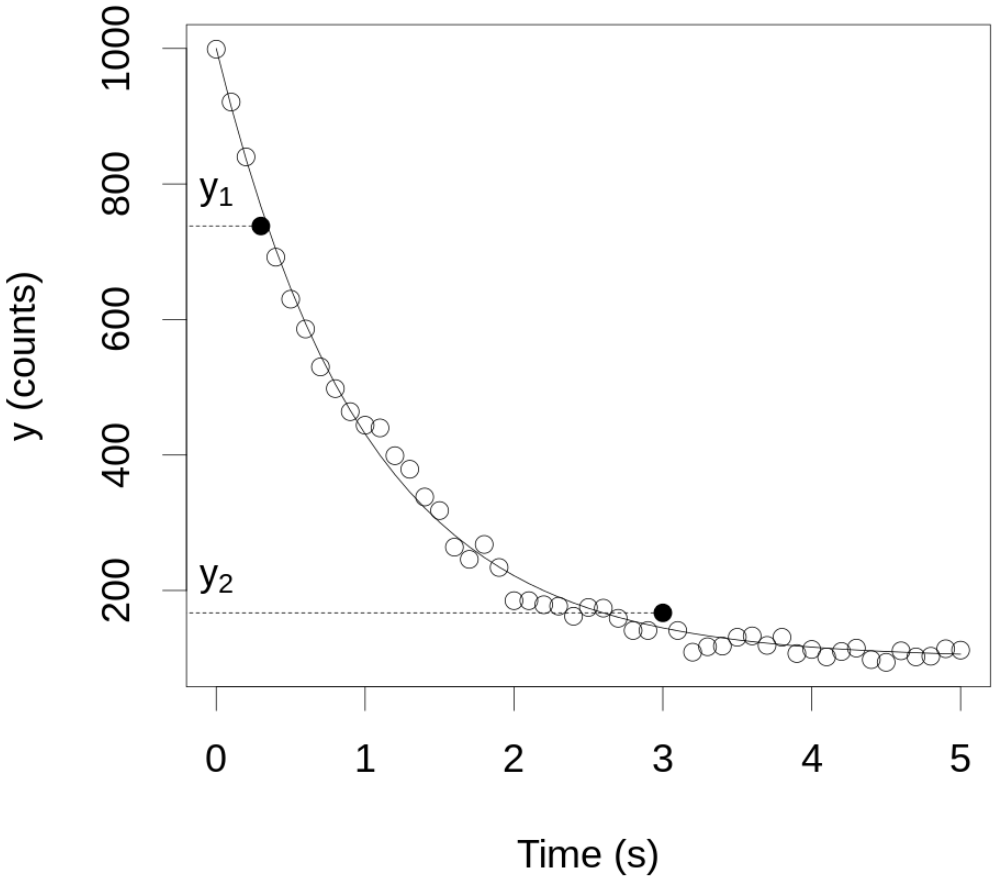

Data simulated according to the previous model. We are going to assume that f_∞ and Δ are known and that (t_1, y_1) and (t_2, y_2) are given. We want to estimate β.

Two estimators (1)

We are going to consider two *estimators* for β:

- The "classical" least square estimator:

$$\tilde{\beta} = \arg\min \tilde{L}(\beta) \ ,$$

where

$$\tilde{L}(\beta) = \sum_j \left(y_j - f(t_j \mid \beta)\right)^2 \ .$$

- The least square estimator applied to the *square root* of the data:

$$\hat{\beta} = \arg\min \hat{L}(\beta) \ ,$$

where

$$\hat{L}(\beta) = \sum_j \left(\sqrt{y_j} - \sqrt{f(t_j \mid \beta)}\right)^2 \ .$$

Two estimators (2)

We perform an empirical study as follows:

- We simulate 100,000 experiments such that:

$$(Y_1, Y_2) \sim \left(\mathcal{P}(f(0.3|\beta_0), \mathcal{P}(f(3|\beta_0)\right) ,$$

  with $\beta_0 = 1$.
- For each simulated pair, $(y_1, y_2)^{[k]}$ $(k = 1, \ldots, 10^5)$, we minimize $\tilde{L}(\beta)$ and $\hat{L}(\beta)$ to obtain: $(\tilde{\beta}^{[k]}, \hat{\beta}^{[k]})$.
- We build histograms for $\tilde{\beta}^{[k]}$ and $\hat{\beta}^{[k]}$ as density estimators of our estimators.

Two estimators (3)

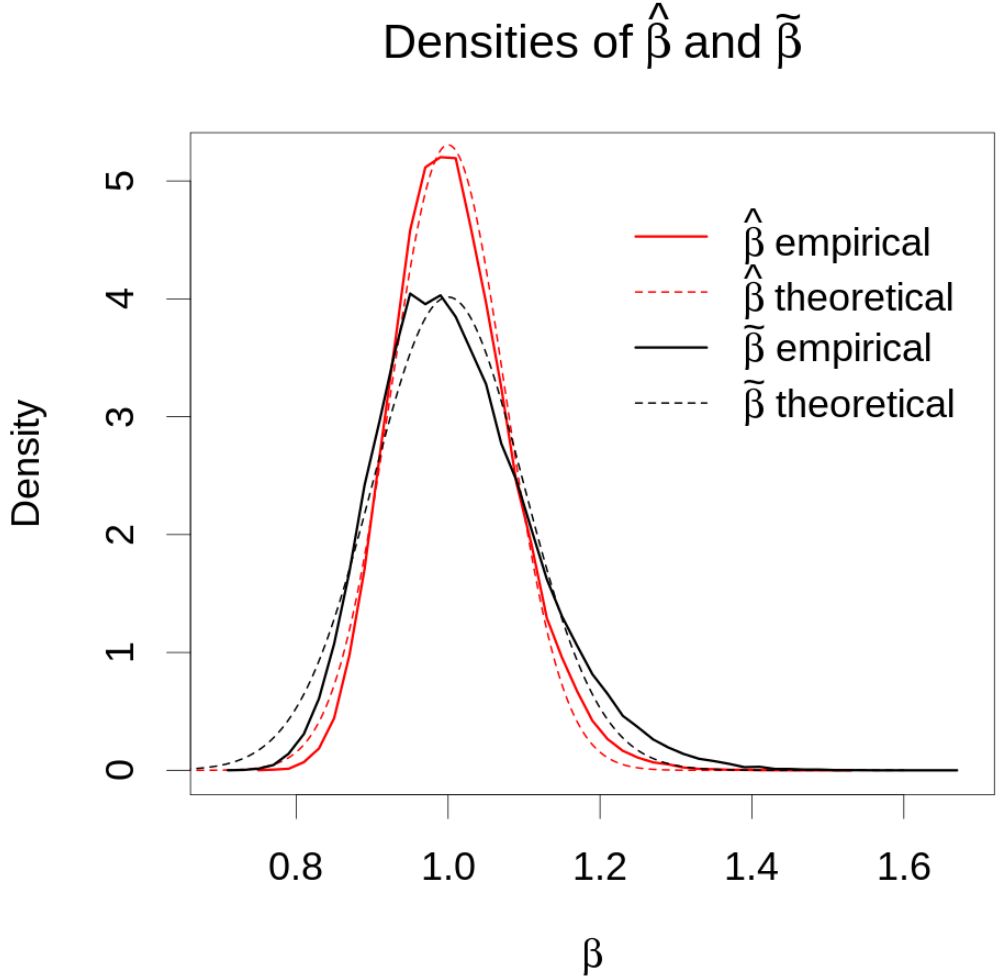

Both histograms are built with 50 bins. $\hat{\beta}$ is clearly better than $\tilde{\beta}$ since its variance is smaller. The derivation of the theoretical (large sample) densities is given in Joucla et al (2010).

CCD basics

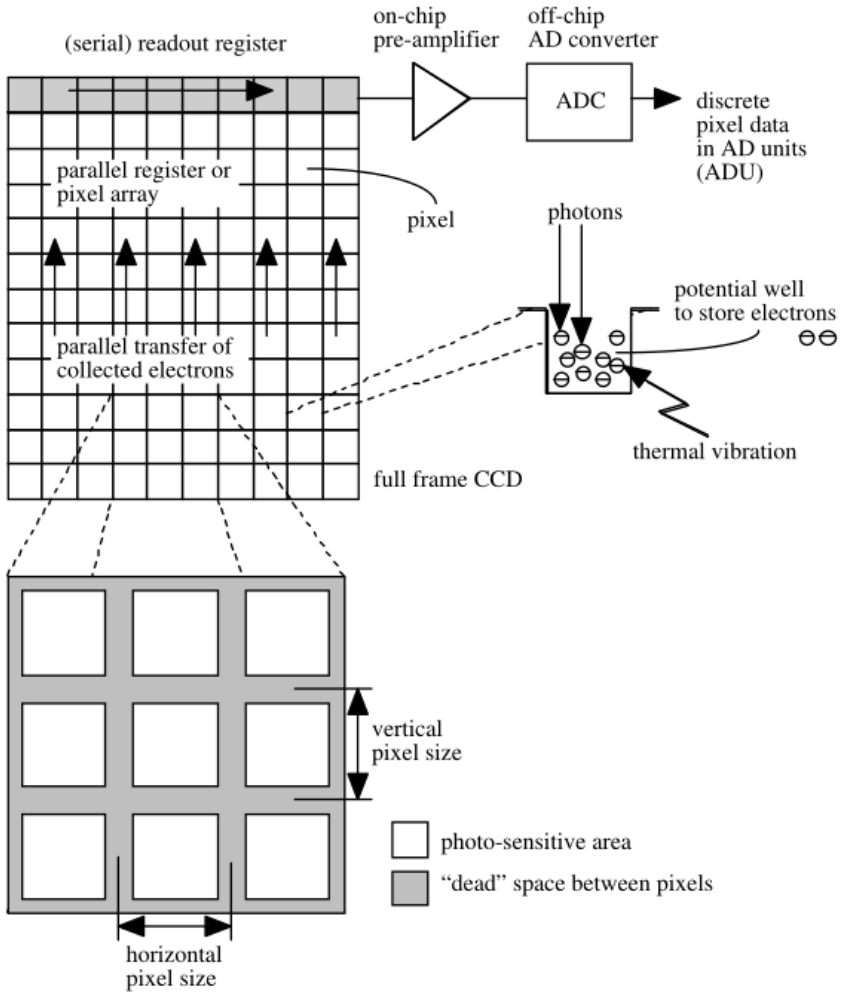

Source: L. van Vliet et col. (1998) Digital Fluorescence Imaging Using Cooled CCD Array Cameras (figure 3).

"Noise" sources in CCD (1)

- The "Photon noise" or "shot noise" arises from the fact the measuring a fluorescence intensity, λ, implies counting photons—unless one changes the laws of Physics there is nothing one can do to eliminate this source of variability (improperly called "noise")—:

$$\Pr\{N = n\} = \frac{\lambda^n}{n!} \exp -\lambda\,, \quad n = 0, 1, \ldots, \quad \lambda > 0\,.$$

- The "thermal noise" arises from thermal agitation which "dumps" electrons in potential wells; this "noise" also follows a Poisson distribution but it can be made negligible by *cooling down* the camera.

"Noise" sources in CCD (2)

- The "read out noise" arises from the conversion of the number of photo-electrons into an equivalent tension; it follows a normal distribution whose variance is independent of the mean (as long as reading is not done at too high a frequency).
- The "digitization noise" arises from the mapping of a continuous value, the tension, onto a grid; it is negligible as soon as more than 8 bit are used.

A simple CCD model (1)

- We can easily obtain a simple CCD model taking into account the two main "noise" sources (photon and read-out).
- To get this model we are going the fact (a theorem) that when a large number of photon are detected, the Poisson distribution is well approximated by (converges in distribution to) a normal distribution with identical mean and variance:

$$\Pr\{N = n\} = \frac{\lambda^n}{n!} \exp -\lambda \approx \mathcal{N}(\lambda, \lambda) \ .$$

- In other words:

$$N \approx \lambda + \sqrt{\lambda}\,\epsilon \ ,$$

where $\epsilon \sim \mathcal{N}(0, 1)$ (follows a standard normal distribution).

A simple CCD model (2)

- A read-out noise is added next following a normal distribution with 0 mean and variance σ_R^2.
- We are therefore adding to the random variable N a new independent random variable $R \sim \mathcal{N}(0, \sigma_R^2)$ giving:

$$M \equiv N + R \approx \lambda + \sqrt{\lambda + \sigma_R^2}\,\epsilon\ ,$$

where the fact that the sum of two independent normal random variables is a normal random variable whose mean is the sum of the mean and whose variance is the sum of the variances has been used.

A simple CCD model (3)

- Since the capacity of the photo-electron weels is finite (35000 for the camera used in the first slides) and since the number of photon-electrons will be digitized on 12 bit (4096 levels), a "gain" G smaller than one must be applied if we want to represent faithfully (without saturation) an almost full well.
- We therefore get:

$$Y \equiv G \cdot M \approx G\,\lambda + \sqrt{G^2\,(\lambda + \sigma_R^2)}\,\epsilon\ .$$

For completeness: Convergence in distribution of a Poisson toward a normal rv (1)

We use the moment-generating function and the following theorem (*e.g.* John Rice, 2007, *Mathematical Statistics and Data Analysis*, Chap. 5, Theorem A):

- If the moment-generating function of each element of the rv sequence X_n is $m_n(t)$,
- if the moment-generating function of the rv X is $m(t)$,
- if $m_n(t) \to m(t)$ when $n \to \infty$ for all $|t| \leq b$ where $b > 0$
- then $X_n \xrightarrow{D} X$.

For completeness: Convergence in distribution of a Poisson toward a normal rv (2)

Lets show that:

$$Y_n = \frac{X_n - n}{\sqrt{n}} \ ,$$

where X_n follows a Poisson distribution with parameter n, converges in distribution towards Z standard normal rv.
We have:

$$m_n(t) \equiv \mathrm{E}\left[\exp(Y_n t)\right] \ ,$$

therefore:

$$m_n(t) = \sum_{k=0}^{\infty} \exp\left(\frac{k-n}{\sqrt{n}} t\right) \frac{n^k}{k!} \exp(-n) \ ,$$

For completeness: Convergence in distribution of a Poisson toward a normal rv (3)

$$m_n(t) = \exp(-n)\exp(-\sqrt{n}t)\sum_{k=0}^{\infty}\frac{\left(n\exp\left(t/\sqrt{n}\right)\right)^k}{k!}$$

$$m_n(t) = \exp\left(-n-\sqrt{n}t+n\exp(t/\sqrt{n})\right)$$

$$m_n(t) = \exp\left(-n-\sqrt{n}t+n\sum_{k=0}^{\infty}\left(\frac{t}{\sqrt{n}}\right)^k\frac{1}{k!}\right)$$

$$m_n(t) = \exp\left(-n-\sqrt{n}t+n+\sqrt{n}t+\frac{t^2}{2}+n\sum_{k=3}^{\infty}\left(\frac{t}{\sqrt{n}}\right)^k\frac{1}{k!}\right)$$

$$m_n(t) = \exp\left(\frac{t^2}{2}+n\sum_{k=3}^{\infty}\left(\frac{t}{\sqrt{n}}\right)^k\frac{1}{k!}\right)$$

For completeness: Convergence in distribution of a Poisson toward a normal rv (4)

We must show:

$$n\sum_{k=3}^{\infty}\left(\frac{t}{\sqrt{n}}\right)^k\frac{1}{k!}\to_{n\to\infty}0\quad\forall\,|t|\le b,\quad\text{where}\quad b>0\,,$$

since $\exp(-t^2/2)$ is the moment-generating function of a standard normal rv. But

$$\left|n\sum_{k=3}^{\infty}\left(\frac{t}{\sqrt{n}}\right)^k\frac{1}{k!}\right|\to_{n\to\infty}0\quad\forall\,|t|\le b,\quad\text{where}\quad b>0$$

implies that since

$$-\left|n\sum_{k=3}^{\infty}\left(\frac{t}{\sqrt{n}}\right)^k\frac{1}{k!}\right|\le n\sum_{k=3}^{\infty}\left(\frac{t}{\sqrt{n}}\right)^k\frac{1}{k!}\le\left|n\sum_{k=3}^{\infty}\left(\frac{t}{\sqrt{n}}\right)^k\frac{1}{k!}\right|.$$

For completeness: Convergence in distribution of a Poisson toward a normal rv (5)

But for all $|t| \leq b$ where $b > 0$

$$\begin{aligned} 0 \leq \left| n \sum_{k=3}^{\infty} \left(\frac{t}{\sqrt{n}} \right)^k \frac{1}{k!} \right| &\leq n \sum_{k=3}^{\infty} \left(\frac{|t|}{\sqrt{n}} \right)^k \frac{1}{k!} \\ &\leq \frac{|t|^3}{\sqrt{n}} \sum_{k=0}^{\infty} \left(\frac{|t|}{\sqrt{n}} \right)^k \frac{1}{(k+3)!} \\ &\leq \frac{|t|^3}{\sqrt{n}} \sum_{k=0}^{\infty} \left(\frac{|t|}{\sqrt{n}} \right)^k \frac{1}{k!} \\ &\leq \frac{|t|^3}{\sqrt{n}} \exp\left(\frac{|t|}{\sqrt{n}} \right) \to_{n \to \infty} 0 \,, \end{aligned}$$

which completes the proof.

For completeness: Convergence in distribution of a Poisson toward a normal rv (6)

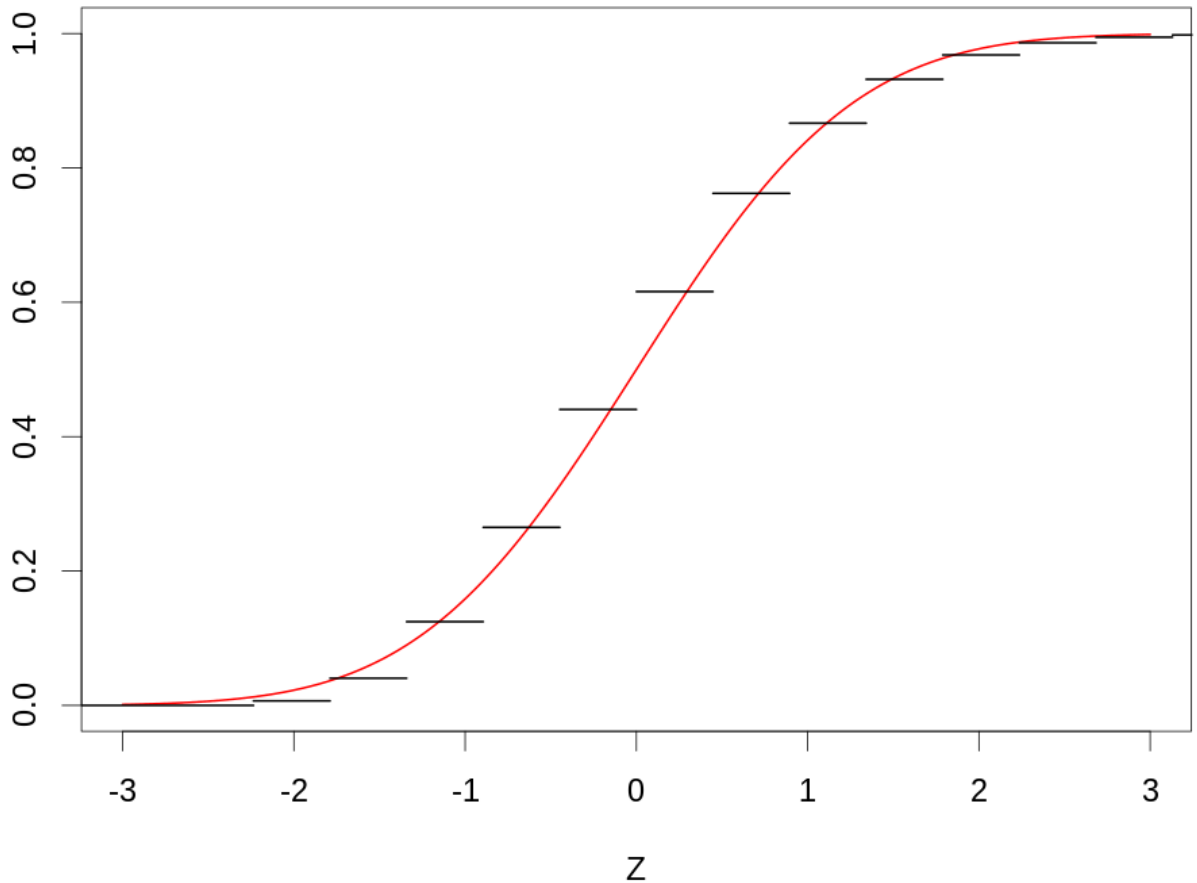

Cumulative distribution functions (CDF) of Y_5 and Z (standard normal).

For completeness: Convergence in distribution of a Poisson toward a normal rv (7)

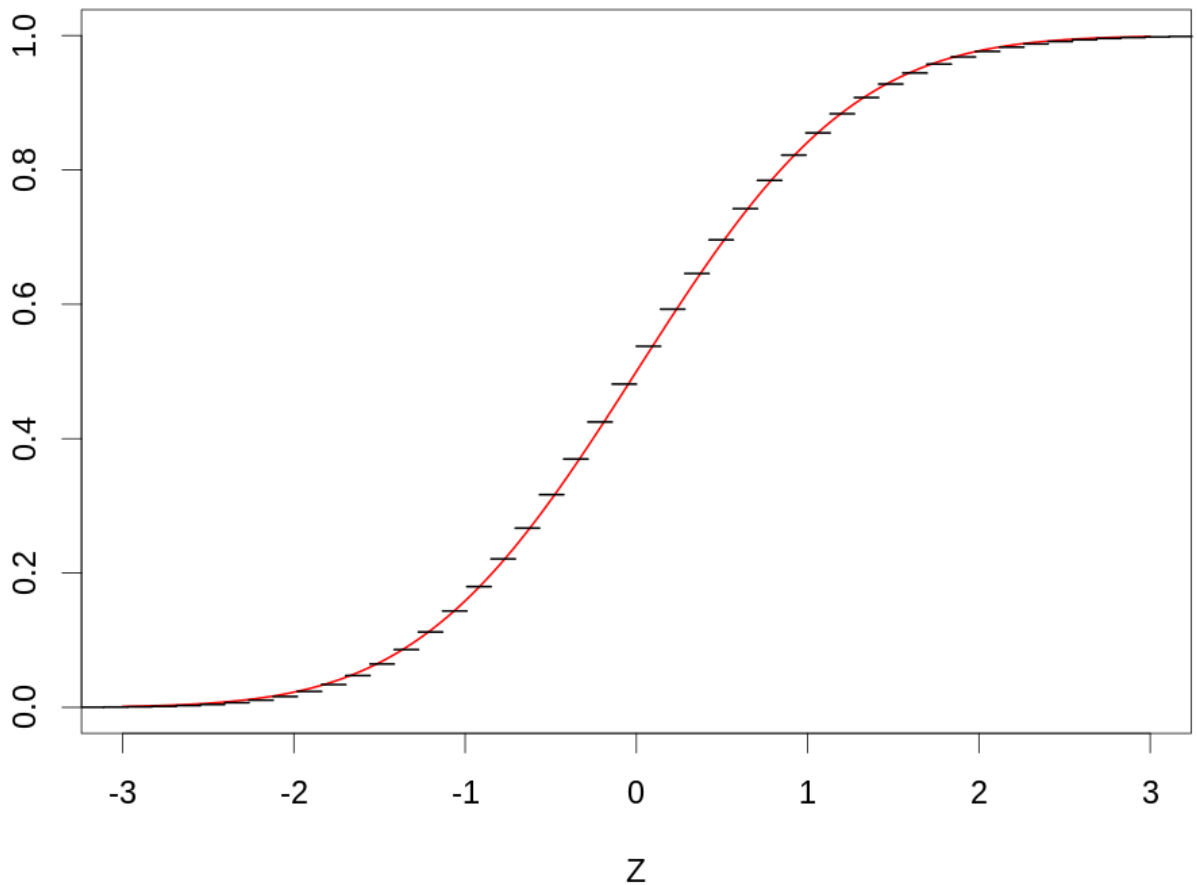

Cumulative distribution functions (CDF) of Y_{50} and Z (standard normal).

For completeness: Convergence in distribution of a Poisson toward a normal rv (8)

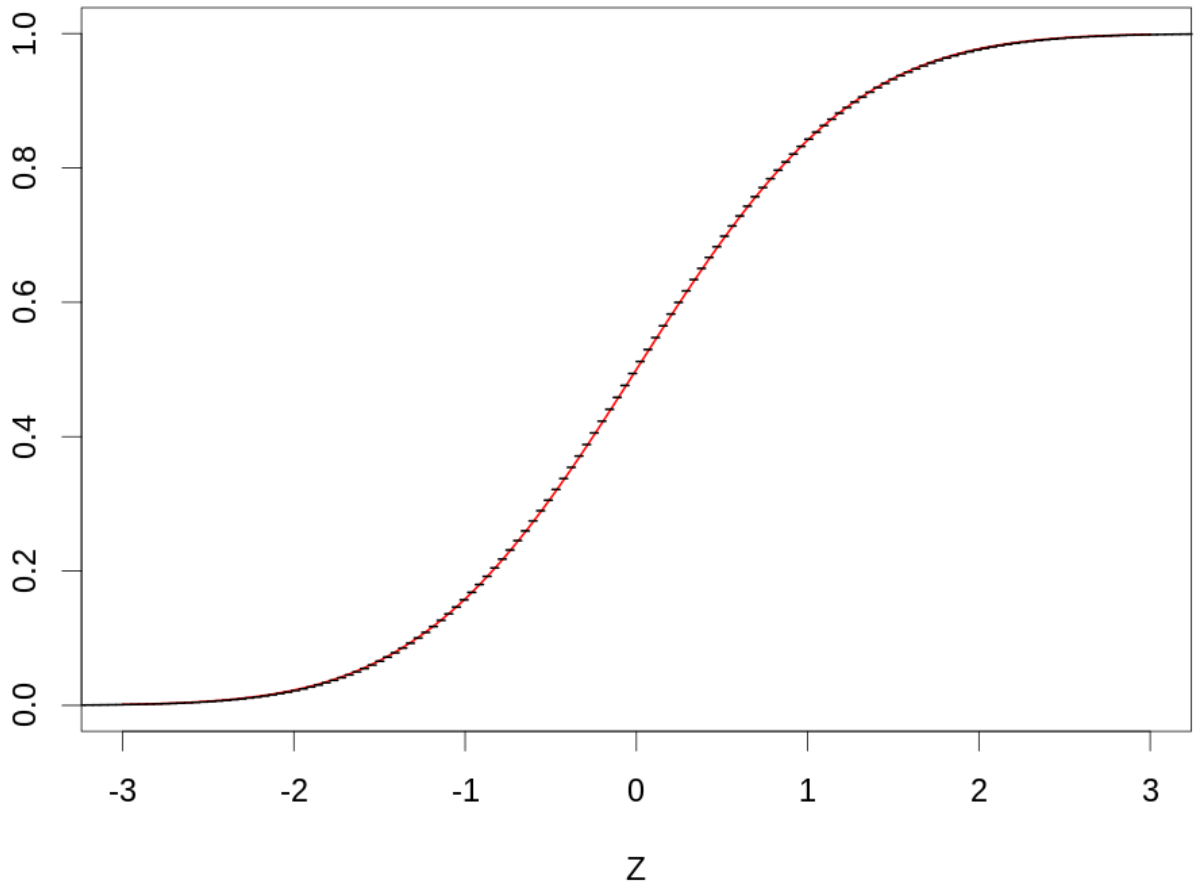

Cumulative distribution functions (CDF) of Y_{500} and Z.

For completeness: Convergence in distribution of a Poisson toward a normal rv (9)

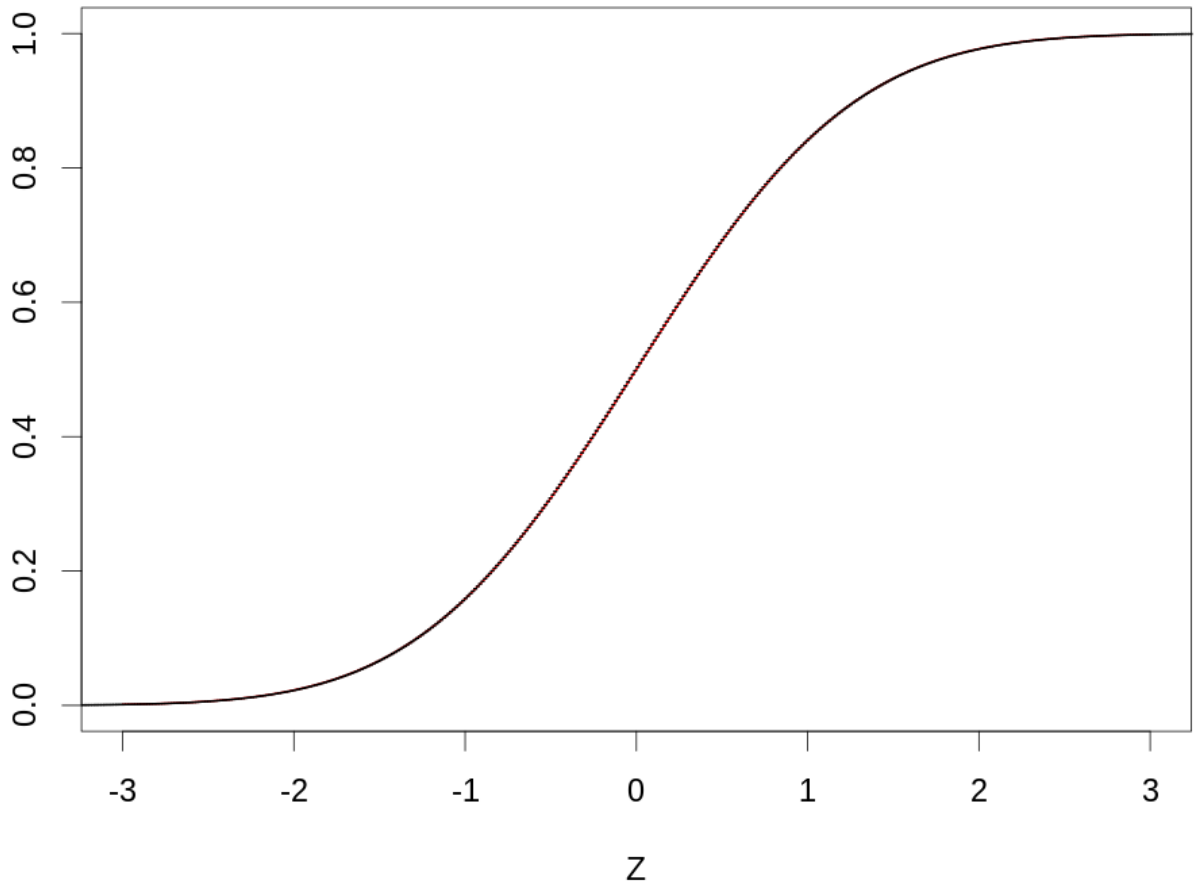

Cumulative distribution functions (CDF) of Y_{5000} and Z.

CCD calibration (1)

If what I just exposed is correct, with the two (main) "noise" sources, the observations Y (from a CCD pixel) follow:

$$Y \sim G\,\lambda + \sqrt{G^2\,(\lambda + \sigma_R^2)}\,\epsilon\ ,$$

where G is the camera gain, σ_R^2 is the read-out variance and ϵ is a standard normal rv. The values of G and σ_R^2 are specified by the manufacturer for each camera, but experience shows that manufacturers tend to be overoptimistic when it comes to their product performances—they can for instance give an underestimated σ_R^2. Its therefore a good idea to measure these parameters with calibration experiments. Such calibration experiments are also the occasion to check that our simple model is relevant.

CCD calibration (2)

- Our problem becomes: How to test $Y \sim G\,\lambda + \sqrt{G^2\,(\lambda + \sigma_R^2)}\,\epsilon$? Or how to set different values for λ?
- Let's consider a pixel of our CCD "looking" at a fixed volume of a fluorescein solution with a given (and stable) concentration. We have two ways of modifying λ :
  - Change the intensity i_e of the light source exciting the fluorophore.
  - Change the exposure time τ.

CCD calibration (3)

We can indeed write our λ as:

$$\lambda = \phi v c i_e \tau \,,$$

where

- v is the solution's volume "seen" by a given pixel,
- c is the fluorophore's concentration,
- ϕ is the quantum yield.

In practice it is easier to vary the exposure time τ and that's what was done in the experiments described next... Question: Can you guess what these experiments are?

CCD calibration (4)

Sebastien Joucla and myself asked our collaborators from the Kloppenburg lab (Cologne University) to:

- choose 10 exposure times,
- for each of the 10 times, perform 100 exposures,
- for each of the 10 x 100 exposures, record the value y_{ij} of the rv Y_{ij} of CCD's pixel i, j.

We introduce a rv Y_{ij} for each pixel because it is very difficult (impossible) to have a uniform intensity (i_e) and a uniform volume (v) and a uniform quantum yield (ϕ). We have therefore for each pixel:

$$Y_{i,j} \sim G\, p_{i,j}\tau + \sqrt{G^2\,(p_{i,j}\tau + \sigma_R^2)}\,\epsilon_{i,j}\ ,$$

where $p_{i,j} = c\phi_{i,j}v_{i,j}i_{e,i,j}$.

CCD calibration (5)

- If our model is correct we should have for each pixel i, j, for a given exposure time, a mean value:

$$\bar{y}_{i,j} = \frac{1}{100} \sum_{k=1}^{1} y_{i,j,k} \approx G\, p_{i,j} \tau$$

- and a variance:

$$S^2_{i,j} = \frac{1}{99} \sum_{k=1}^{1} (y_{i,j,k} - \bar{y}_{i,j})^2 \approx G^2 \left(p_{i,j} \tau + \sigma^2_R \right) .$$

- The graph of $S^2_{i,j}$ *vs* $\bar{y}_{i,j}$ should be a straight line with slope G ordinate at 0, $G^2 \sigma^2_R$.

CCD calibration (6)

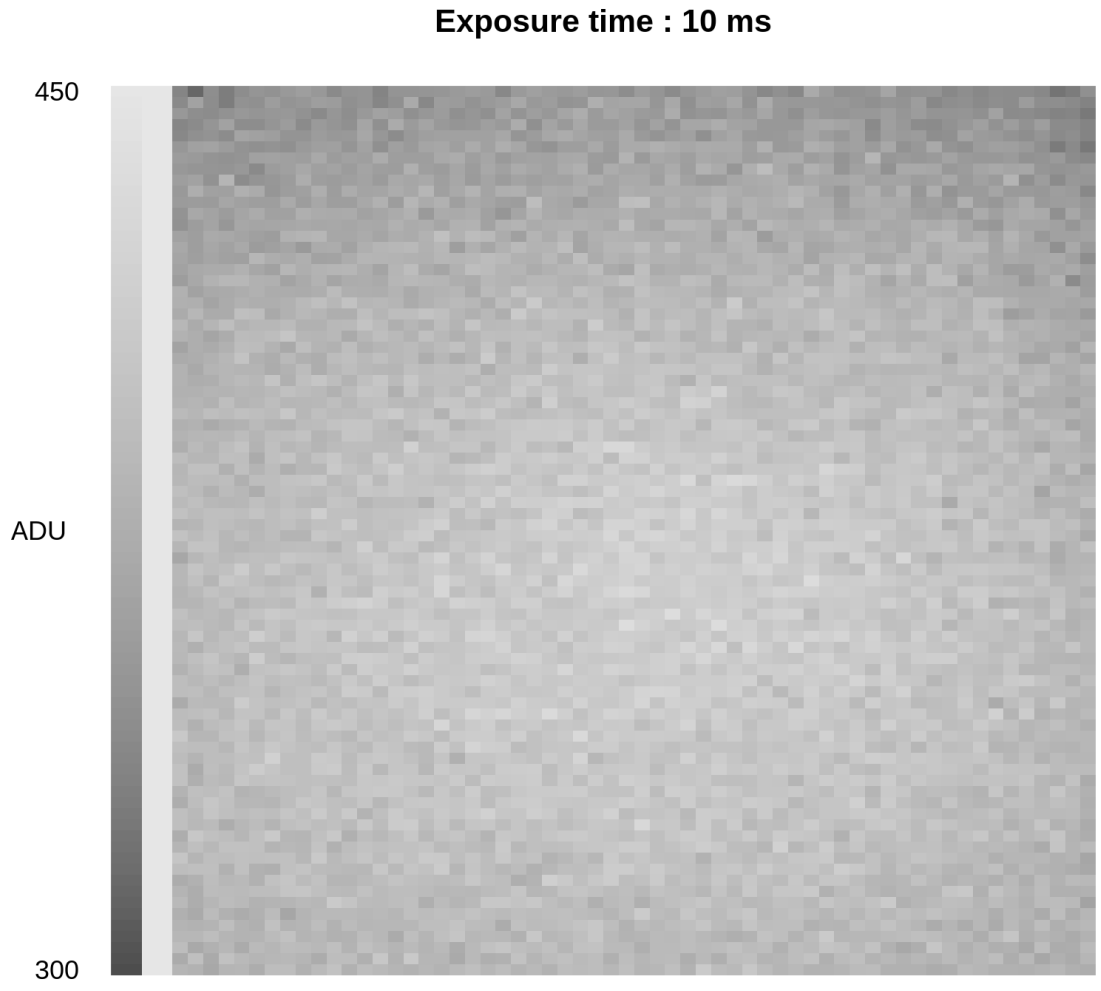

The first exposure of 10 ms (experiment performed by Andreas Pippow, Kloppenburg lag, Cologne University).

CCD calibration (7)

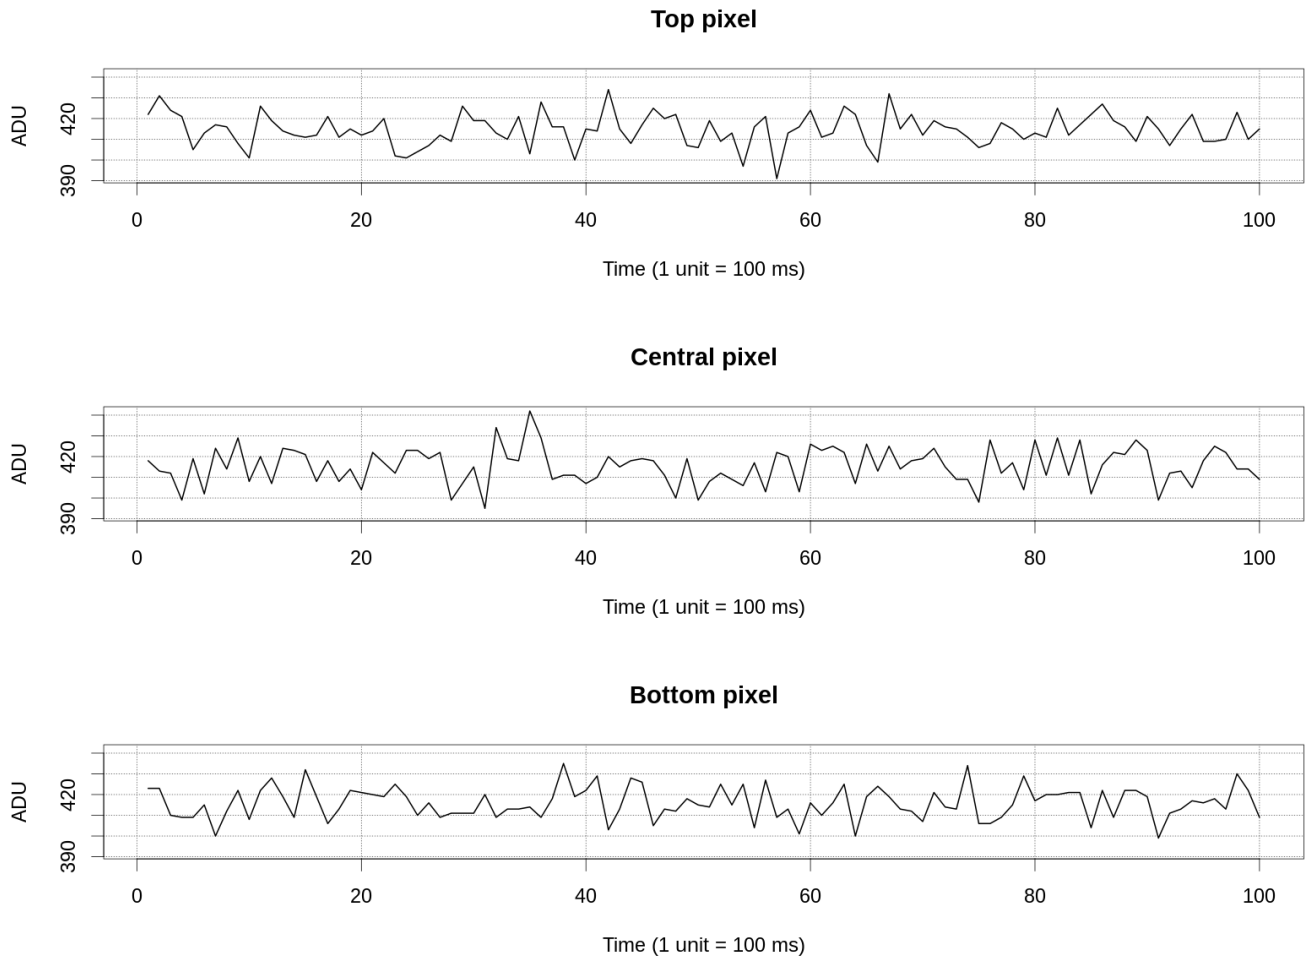

Counts time evolution for three neighboring pixels (10 ms exposure time).

CCD calibration (8)

- The data are going to be analyzed as if the $Y_{i,j,k}$ were IID, but they were sequentially recorded. It is therefore strongly recommended to check that the IID hypothesis is reasonable.
- The small example of the previous figure shows that there are no (obvious) trends.
- We must also check the correlation function.

CCD calibration (9): absence of correlations

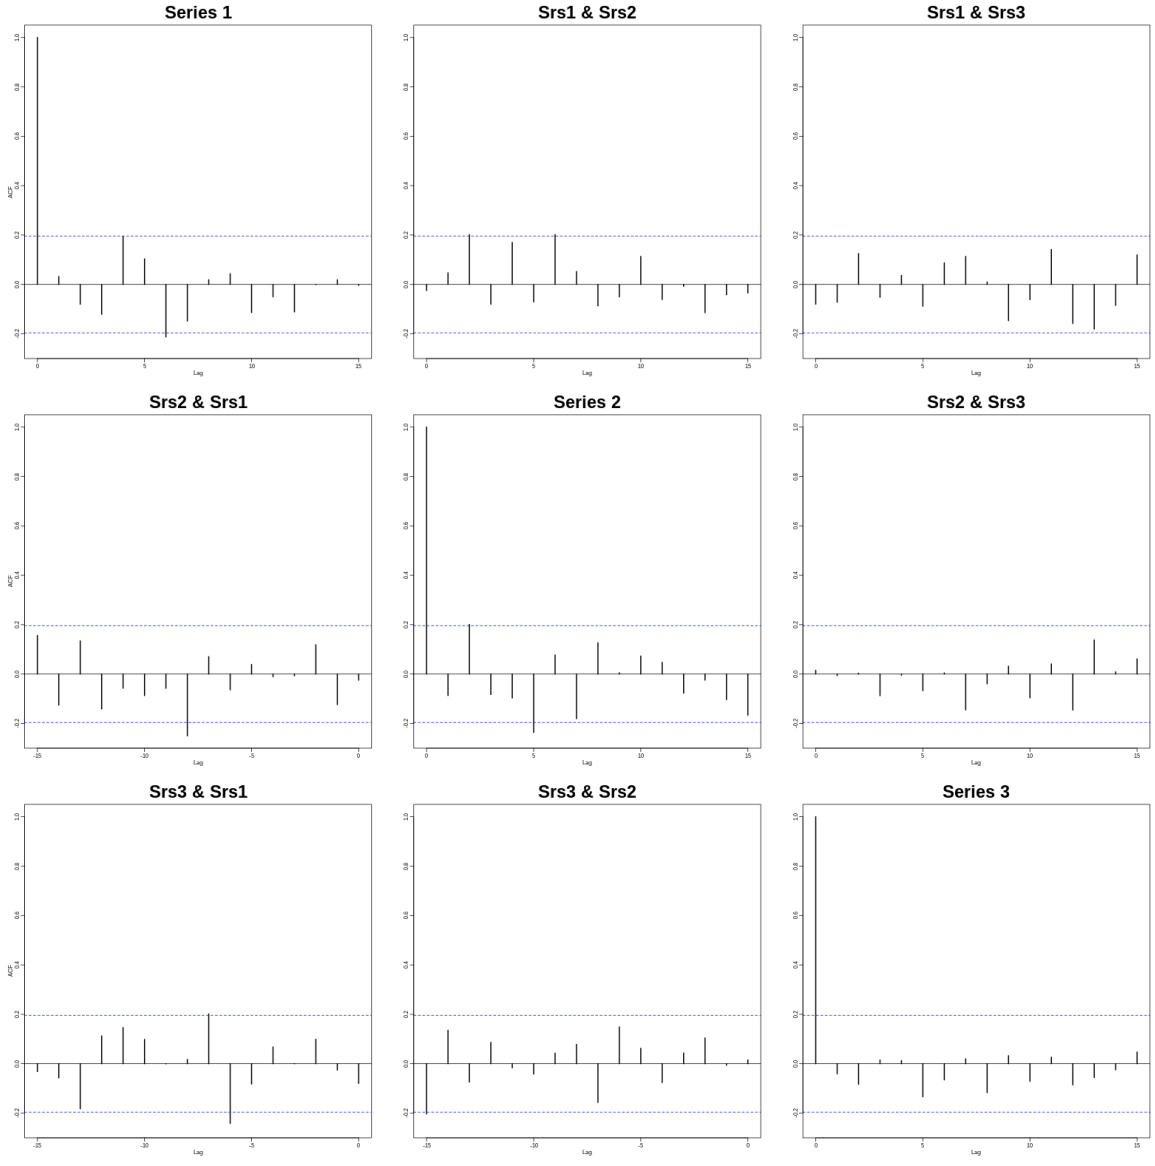

CCD calibration (10): $S^2_{i,j}$ vs $\bar{y}_{i,j}$

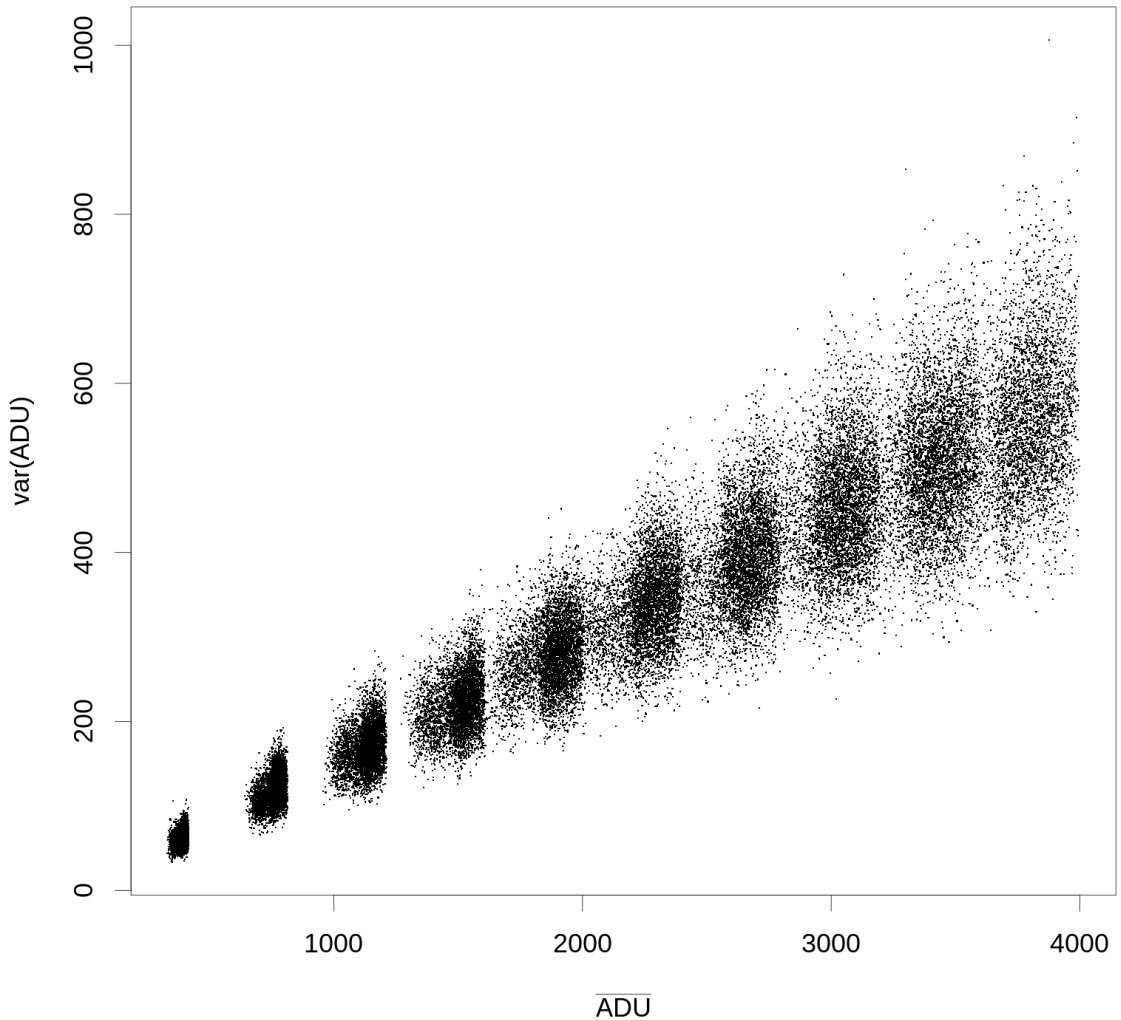

We do see the expected linear relation:
$\mathrm{Var}[ADU] = G\mathrm{E}[ADU] + G^2\sigma^2_R$.

CCD calibration (11): Linear fit

The heteroscedasticity (inhomogeneous variance) visible on the graph is also expected since the variance of a variance for an IID sample of size n from a normal distribution with mean μ and variance σ^2 is:

$$\mathrm{Var}[S^2] = \frac{2\sigma^4}{(n-1)} \ .$$

- This means than when we do our linear fit we should use weights.
- A software package like R allows us to do that by giving a vector whose elements are proportional to inverse of the variance.

CCD calibration (12): Linear fit

- If our mean values are in a variable called `ADU.m` and our variances in a variable called `ADU.V`, we would use the following call:

```
varVSmean <- lm(ADU.v ~ ADU.m, weights = 99/2/ADU.v^2)
```

- And get:

```
(Intercept)       ADU.m
      5.611       0.140
```

CCD calibration (13): fit checking

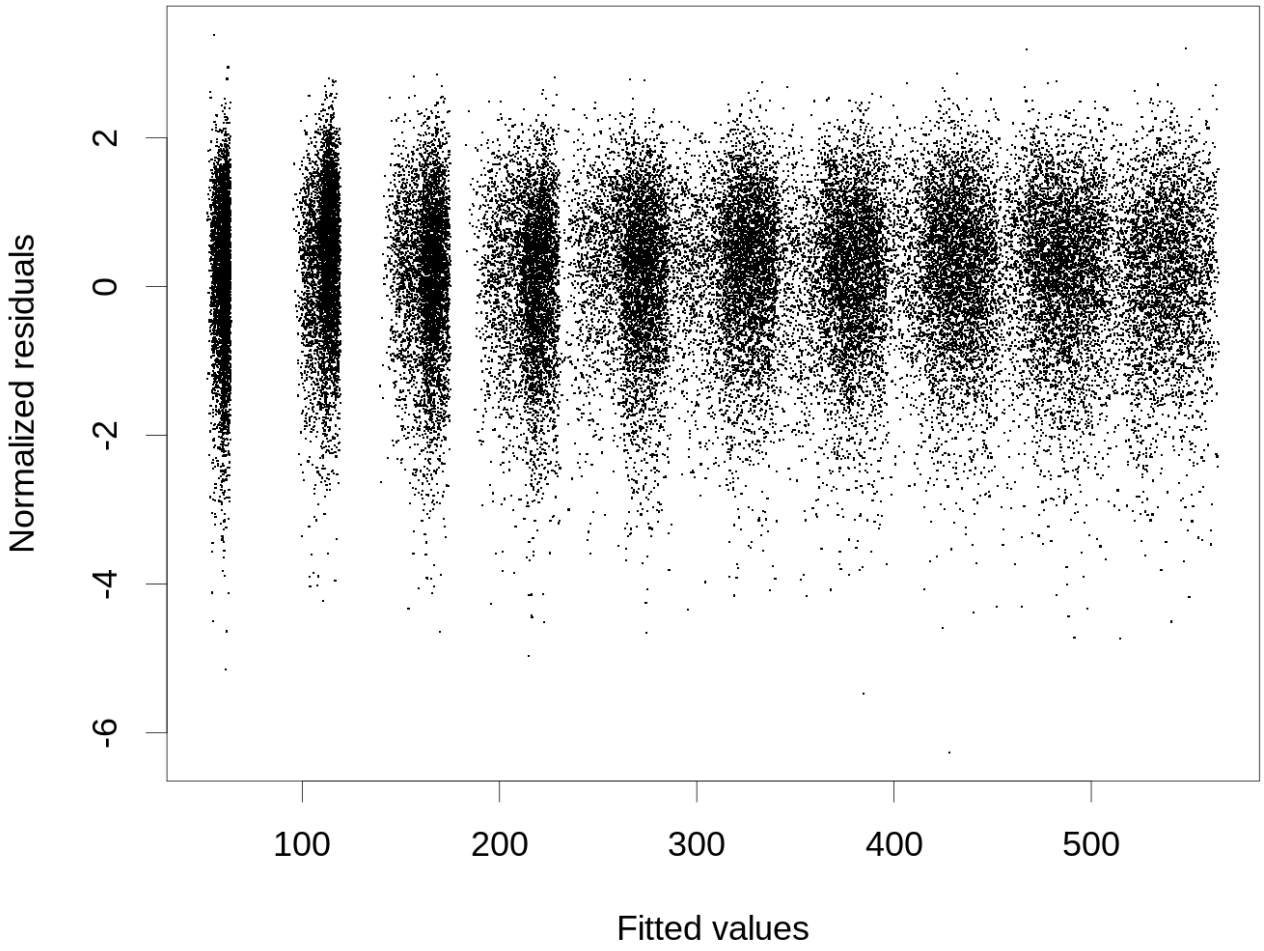

CCD calibration (14): some remarks

- When we use a linear regression, we are (implicitly) assuming that the "independent" variable, here `ADU.m`, is *exactly* known.
- This was clearly not the case here since `ADU.m` was measured (with an error).
- We could therefore refine our fit.

Error propagation

- Let us consider three random variables: X, Y and Z such that:
- $X \approx \mathcal{N}(\mu_X, \sigma_X^2)$ or $X \approx \mu_X + \sigma_X \, \epsilon$ $(\epsilon \sim \mathcal{N}(0,1))$
- $Y \approx \mathcal{N}(\mu_Y, \sigma_Y^2)$ or $Y \approx \mu_Y + \sigma_Y \, \epsilon$
- X and Y are independent
- $Z = f(X, Y)$, with f continuous and differentiable.
- Using a first order Taylor expansion we then have:

$$\begin{aligned} Z &\approx f(\mu_X + \sigma_X \, \epsilon_1, \mu_Y + \sigma_Y \, \epsilon_2) \\ &\approx f(\mu_X, \mu_Y) + \sigma_X \, \epsilon_1 \, \frac{\partial f}{\partial X}(\mu_X, \mu_Y) + \sigma_Y \, \epsilon_2 \, \frac{\partial f}{\partial Y}(\mu_X, \mu_Y) \end{aligned}$$

- $\mathrm{E}Z \approx f(\mu_X, \mu_Y) = f(\mathrm{E}X, \mathrm{E}Y)$
- $\mathrm{Var}Z \equiv \mathrm{E}[(Z - \mathrm{E}Z)^2] \approx \sigma_X^2 \, \frac{\partial f}{\partial X}^2(\mu_X, \mu_Y) + \sigma_Y^2 \, \frac{\partial f}{\partial Y}^2(\mu_X, \mu_Y)$
- $Z \approx f(\mu_X, \mu_Y) + \sqrt{\sigma_X^2 \, \frac{\partial f}{\partial X}^2(\mu_X, \mu_Y) + \sigma_Y^2 \, \frac{\partial f}{\partial Y}^2(\mu_X, \mu_Y)} \, \epsilon$

Variance stabilization (1): Theory

- For our CCD model we have (for a given pixel):

$$Y \sim G\,\lambda + \sqrt{G^2\,(\lambda + \sigma_R^2)}\,\epsilon = \mu_Y + \sqrt{G\,\mu_Y + G^2\,\sigma_R^2}$$

- Then if $Z = f(Y)$ we get:

$$Z \approx f(\mu_Y) + \mid f'(\mu_Y) \mid G\sqrt{\mu_Y/G + \sigma_R^2}\,\epsilon$$

- What happens then if we take: $f(x) = 2\sqrt{x/G + \sigma_R^2}$?
- We have:

$$f'(x) = \frac{1}{G\sqrt{x/G + \sigma_R^2}}$$

- Leading to:

$$Z \approx 2\sqrt{\mu_Y/G + \sigma_R^2} + \epsilon$$

Variance stabilization (2): Example

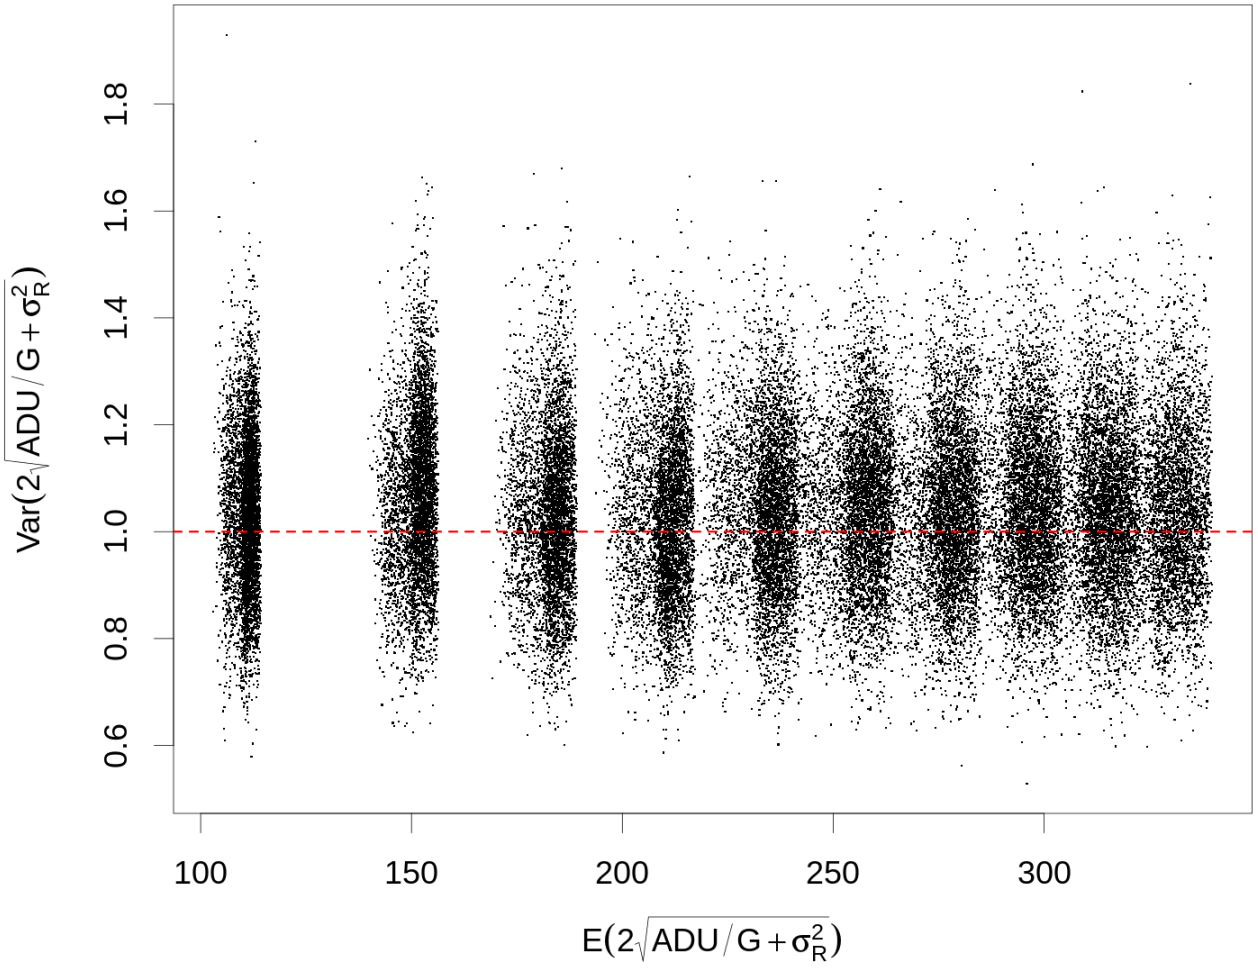

Back to where we started

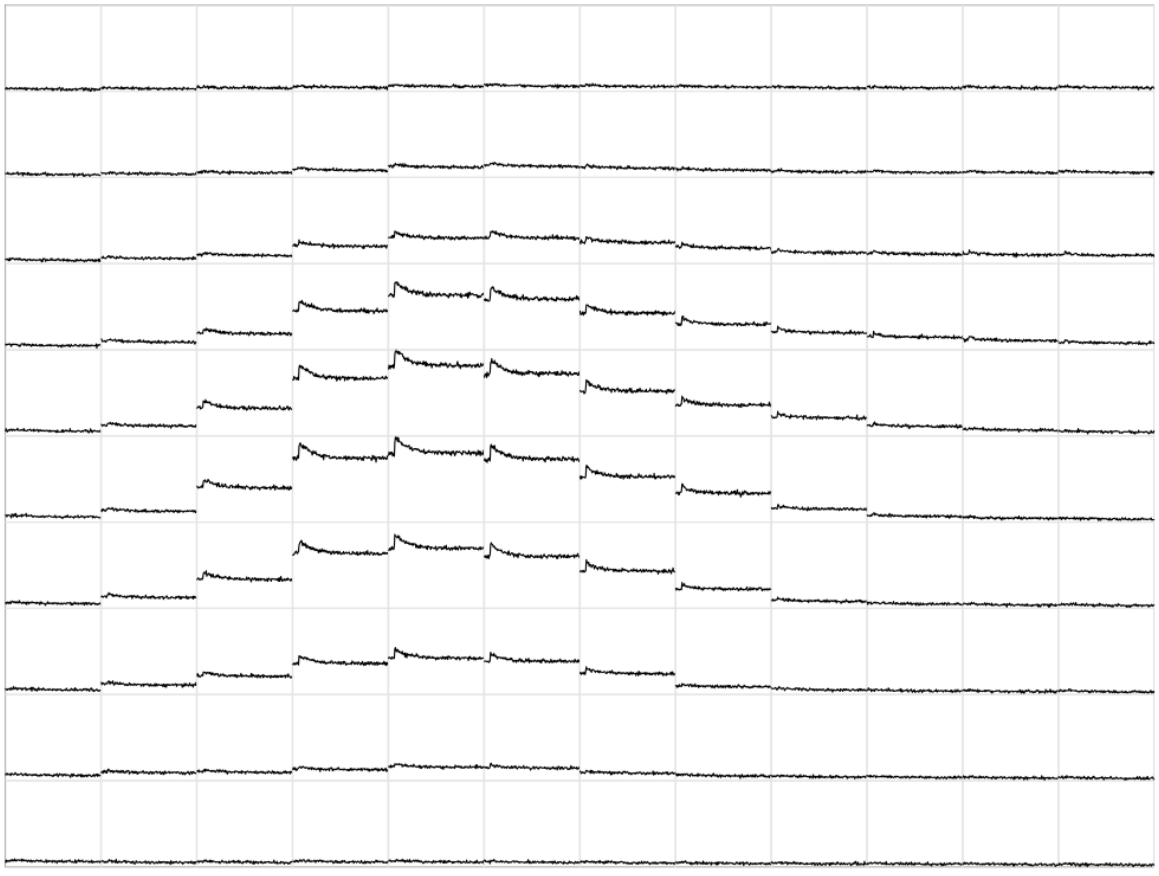

ADU counts (raw data) from Fura-2 excited at 340 nm. Each square corresponds to a pixel. 25.05 s of data are shown. Same scale on each sub-plot. 12 x 10 among 60 x 80 pixels are shown. Data recorded by Andreas Pippow (Kloppenburg Lab. Cologne University).

Quick ROI detection (1): Motivation

- After variance stabilization: $Z_{i,j,k} = 2\sqrt{ADU_{i,j}/G + \sigma_R^2}$, the variance at each pixel (i, j) at each time, k, should be 1.
- If a pixel contains no dynamical signal—that is nothing more than a constant background signal—the following statistics:

$$RSS_{i,j} \equiv \sum_{k=1}^{K} (Z_{i,j,k} - \overline{Z}_{i,j})^2 \quad \text{with} \quad \overline{Z}_{i,j} \equiv \frac{1}{K} \sum_{k=1}^{K} Z_{i,j,k}$$

  should follow a χ^2 distribution with $K - 1$ degrees of freedom.
- We could therefore compute the values of the complementary cumulative distribution function of the theoretical χ^2_{K-1} distribution:

$$1 - F_{\chi^2_{K-1}}(RSS_{i,j})$$

  and look for very small values—that is very small probabilities—(using a log scale helps here).

Quick ROI detection (2)

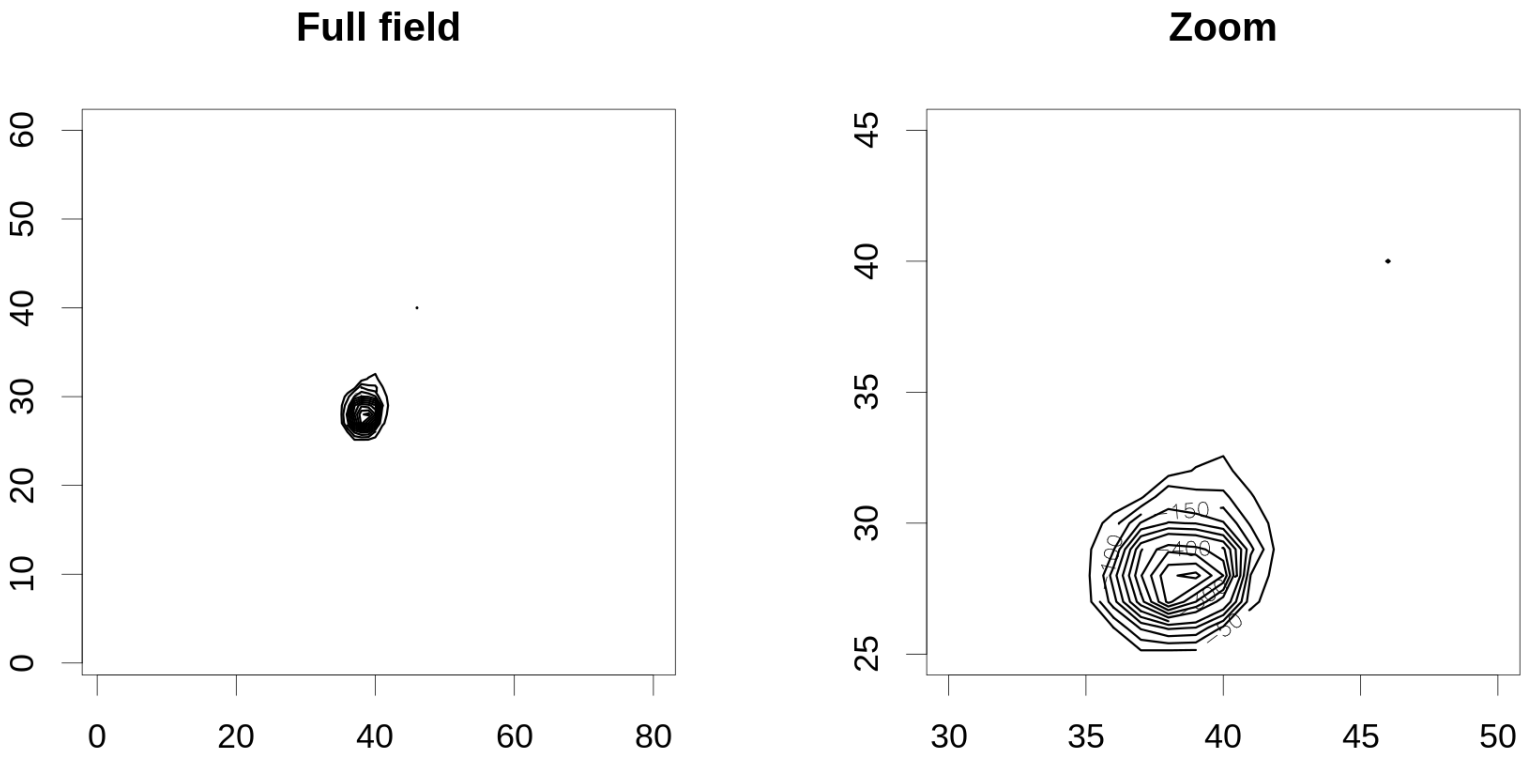

Contour plots of $\log\left(1 - F_{\chi^2_{K-1}}(RSS_{i,j})\right)$

Pointwise time course estimation (1)

- We are going to be (very) conservative and keep as our ROI the pixels having an $\log\left(1 - F_{\chi^2_{K-1}}(RSS)\right) \leq -300$.
- We are then left with 12 pixels.
- We are going to model the fluorescence intensity of each of these pixels by:

$$S_{i,j}(t) = \phi_{i,j}\, f(t) + b\ ,$$

  where $f(t)$ is a signal time course to all pixels of the ROI, $\phi_{i,j}$ is a pixel specific parameter and b is a background fluorescence assumed identical for each pixel.
- The time t is in fact a discrete variable, $t = \delta\, k$ ($\delta = 150$ ms) and we are seeking a pointwise estimation: $\{f_1, f_2, \ldots, f_K\}$ ($K = 168$) where $f_k = f(\delta\, k)$.
- We end up with 12 ($\phi_{i,j}$) + 168 (f_k) + 1 (b) = 181 parameters for 12 x 168 = 2016 measurements.

Pointwise time course estimation (2)

- We need to add a constraint since with our model specification:

$$S_{i,j,k} = \phi_{i,j} f_k + b \ ,$$

  we can multiply all the $\phi_{i,j}$ by 2 and divide all the f_k by 2 and get the same prediction.
- We are going to set $f_1 = 1$ and our pointwise estimation relates to what is usually done with this type of data, $\Delta S(t)/S(1)$ through:

$$\Delta S(t)/S(1) = \frac{S(t) - S(1)}{S(1)} = f(t) - 1 + \text{noise}$$

  .
- Notice that no independent background measurement is used.

Pointwise time course estimation (3)

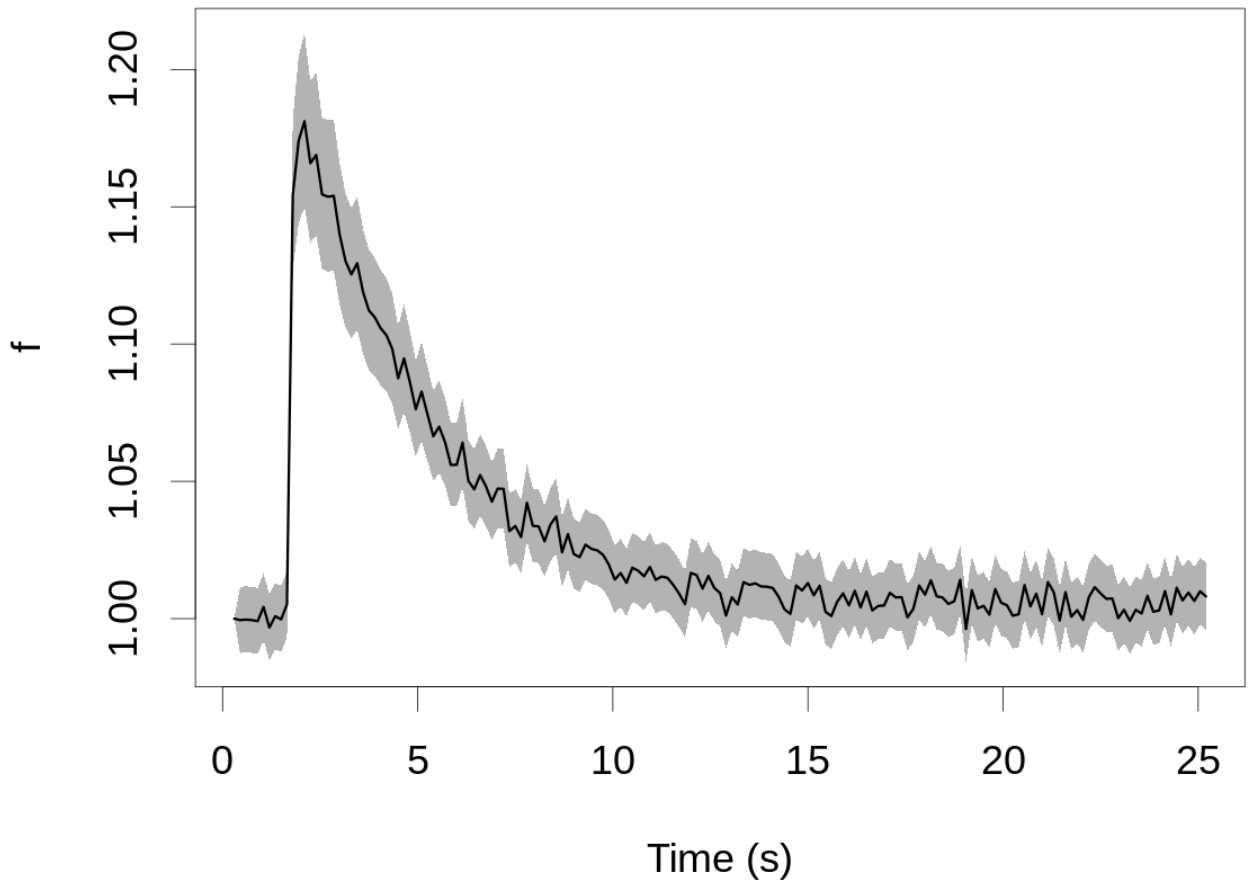

Notice the confidence intervals

Pointwise time course estimation (4)

- In addition to confidence intervals on model parameters, this approach gives classical χ^2 based goodness-of-fit tests.
- Here we get a slightly too large p value: 0.968
- This is likely due to some background inhomogeneity.
- To learn how to deal with that, check: Joucla et al (2013) Estimating background-subtracted fluorescence transients in calcium imaging experiments: A quantitative approach. *Cell Calcium* in press.

Warning

- I haven't considered diffusion effects here, but this is more a model issue than a noise issue.
- I assumed that talking about concentration was meaningful. This is fine when we look at large neurites. Since 1 μM gives roughly 600 ions per cubic μm, a baseline concentration of 50 nM gives 30000 ions is a small soma (10 μm in radius) but gives only 30 ions in a spine.
- In such small volumes, the discrete nature of the ions cannot be ignored anymore, but that's another story...

Thanks

This work was done in collaboration with:

- Sebastien Joucla
- Romain Franconville
- Andeas Pippow
- Peter Kloppenburg

Thank you for your attention!