MEETING OF THE ROYAL ASTRONOMICAL SOCIETY

Friday 2025 March 14 at $16^{h}\ 00^{m}$
in the Geological Society Lecture Theatre, Burlington House

Mike Lockwood, *President*
in the Chair

The President. Welcome to this hybrid meeting. Questions can be put in the Q & A section and they will be read out by Dr. Pam Rowell. Philip Diamond is retiring at this year's AGM after seven years with us, and we thank him for all he has done. He will be replaced by Ian Russell, who will take over on May 1st. The Caroline Herschel Prize Lecture was established by the Herschel Society and it is awarded to UK and German women researchers in alternate years by the RAS and its German equivalent. More details can be found on the RAS website. Unfortunately the current holder, Professor Stefanie Walch-Gassner, who was due to address us, was taken ill at short notice and is unable to be here. Instead I will give a talk entitled 'The Great Aurorae of 2024'.

Last year was a truly remarkable one for aurorae. Aurorae are caused by energetic particles, mostly electrons, hitting our upper atmosphere then precipitating down field lines, hitting atomic oxygen atoms and emitting light. There is a wide range of colours produced: the green appears lower down whilst collisions higher in the atmosphere, due to lower-energy electrons, produce red light. Excitation of molecular nitrogen produces blue light.

Magnetic reconnection happens in the upper atmosphere on the dayside, allowing low-energy electrons to come in down the field lines, producing red aurorae in both hemispheres. Particles can also bounce in the converging magnetic-field lines and end up in the tail plasma sheet; they are accelerated using energy from the solar wind and they give rise to the green aurorae. However, it is particles that come up from the ionosphere that cause most of the aurorae. In quiet periods it is the solar-wind particles that dominate but in disturbed times the major role is taken by the ionosphere, but the energy that comes from the solar wind does always dominate.

Aurorae form an oval around each pole which expands in radius and width, being most intense at the North Cape of Norway, but it can extend to the UK and further south during periods of intense activity. Aurorae often form arches or arcs, if seen from the side, but at lower latitudes the displays are very diffuse and red in colour.

We have conducted a survey of auroral sightings, of which we have 224 870 records going back to 1650, and it appears that most recent sightings are land-based whereas it used to be transatlantic trade and shipping that produced most reports. When you go to Alaska, Iceland, or northern Scandinavia it is very much a cultural activity to go out and observe and record aurorae. At lower latitudes there is intense activity in reporting aurorae and the population density is higher so it is not surprising that auroral sightings correlate with population.

In 2024 there were 4735 reports of aurorae. There was an intense event on May 10/11 and there were four aurorae seen far from the pole in that calendar year. The May 10 event was caused by a coronal mass ejection (CME). It was observed and there was a prediction made about six hours before the first event but, in fact, there were seven CMEs which pummelled the Earth.

It is very tempting to compare the events of 2024 with the Carrington Event of

1859 which was the first to affect directly technology on Earth. The geomagnetic deflection in 2024 May amounted to 400 nT whereas the Carrington event is thought to have reached about 1600 nT. The May 10/11 event is the third lowest latitude reached since records began. In Ontario the aurora was entirely green whereas in Silbury Hill, Wiltshire, we saw red on green with a bit of nitrogen blue in the middle, and in Romania it appeared red over green but here the green formed a narrow band. In Arizona the aurora was almost entirely pure red with a bit of blue. It was also visible from Santa Cruz de Tenerife but the lowest latitude where it was detected was Oman.

The field lines connect to the plasma sheet and the energetic (1 keV) electrons come down the field lines to give a green aurora, then we get energetic particles moving in during the geomagnetic storm from this outer radiation belt and they tip particles that are already there in a low-energy torus called the plasmasphere down onto the Earth and this gives rise to the red colour. High-latitude aurorae are very structured but low-latitude events are very diffuse but more colourful.

In history there are several verses of the *Bible* which are almost certainly descriptions of aurorae. The most interesting is in the second book of Ezekiel where he describes it in such detail and such length that it is almost certainly a red coronal aurora. We think he was living in Nippur, an archaeological site in Iraq. It tallies with Assyrian and Babylonian documents which date it to 567 BC March 12.

Early observers were doctors, clergymen, affluent gentlemen scientists, lighthouse keepers on solitary nights, professional observers, and more recently the latest robotic spacecraft and the public, through citizen science and the social media.

The magnetic pole has moved over time and has affected the appearance of aurorae over the mid-Atlantic and North Africa, neither of which are great places for auroral records. In the 17th Century, Plymouth was a good place to see aurorae but the events slowly went to higher and higher geomagnetic latitudes, whereas the Faroe Islands have moved into the auroral oval and so it is a much better place to see them. Today we have a lot of social media and the modern mobile phone has a great camera for recording aurorae. The eye is as good as a digital camera; if left in the dark for two hours it can detect individual photons, but normally people are not dark adapted.

In terms of the most significant displays, the most intense event was at Khartoum in 1872, followed by one at Panama City in 1859 and one in Oman in 2024, then the Carrington Event of 1859, followed by two events from 2024.

There is some fantastic science to come out. There were two new radiation belts introduced by the May event which was completely unexpected and which are still there.

Professor Mark Lester. Thank you very much, Mike, for speaking to us at such short notice. Are there any questions?

Professor Kathy Whaler. Thanks Mike, really interesting. Do we actually know the magnetic latitude of the site in Iraq in Ezekiel's time?

The President. Will Brown at the BGS did the calculations for me. He used IGRF (International Geomagnetic Reference Field) and then went further back in time by splining it with gufm1 [the model describing the magnetic field at the core–mantle boundary for 1590–1990] and then we used something called a modified quasi-dipole.

Professor Whaler. So you don't actually know it? You are just inferring it as records don't go beyond 1590 or so with gufm1. So, you are extrapolating 1200 years?

The President. We have the IGRF values for the last century but before 1900 we have to use gufm1 and then extrapolate to biblical times.

Mr. Horace Regnart. Thank you. Speaking as a Northumbrian, people might like to know that the first of your photos showed our national tree, the sycamore, against a wonderful aurora. It is now coming back to life. Isn't that wonderful?

The President. It is! That is the best news I've had today. I don't know quite why that upset me as much as the sheer wanton stupidity of it.

Ms. Ana Vitiello. You didn't mention anything about the Southern Hemisphere. Do you have any data on that please?

The President. There are data. The reason we didn't use them in our survey is for two reasons: firstly, an awful lot of the southern hemisphere is water which means that the possibilities are much lower, but, historically, after about 1650, Europe really woke up to measuring around the world and got quite good statistics then, but we don't have any equivalent in the southern hemisphere and so the gradient of the number of potential observers is quite different. And so we decided best just to stick to one hemisphere. There are data and they are quite interesting sometimes.

Mr. Samir Chitnavis. Thank you for a very colourful talk on a Friday afternoon. I research photosynthesis under different coloured stars, which means I get to grow algae under rainbow-coloured lights. I wanted to ask you, are there any observations of animals changing their behaviour during an aurora? Other than humans!

The President. If there are, I would really like to know about it! You raise a very interesting point. One of the big problems now is we get a lot of reports of red aurorae that are actually greenhouses because they use a mixture of red and blue LEDs to make things grow faster, and if it's cloudy it looks like aurorae, and sports stadia where somebody's posted a picture of the aurora at Stamford Bridge. It's where they are using a mixture of lights to get the grass to grow back faster and faster, and so that is an increasing problem of light pollution.

Professor Lester. Thank you again, Mike [applause].

The President. Our next speaker is Andy Smith who is going to tell us about space weather with AI, so he is going to pick up on the sort of applications we are talking about. He was awarded his PhD from Southampton in 2018, and his thesis focussed on physics and plasma environments around Mercury and primarily space weather. He has returned to Earth in every sense and worked at UCL where he combined physics expertise, AI, and interrogative space-weather forecasting methods. In 2022 he was awarded a NERC Independent Research Fellowship which he currently holds at Northumbria University. The title of his talk is 'Space-weather forecasting, geomagnetically induced currents & machine learning'.

Dr. Andy Smith. The take-home message from this talk is that we don't need complex models, we need science.

The aurora is a fantastic manifestation of the interaction of the outflowing solar wind with our geomagnetic environment. The aurora is a manifestation of the ionospheric electrical current that flows in the atmosphere and induces electromagnetic fields and electric fields, and over the past 50 to 100 years we have been putting longer and longer conductors across the Earth which provide very nice avenues to bridge over these geoelectric fields and consequently place geomagnetically induced currents (GICs) in power networks and pipelines. These happen every day but we really feel them when they get really large. They add what is primarily a DC offset into the AC systems — leading to direct damage, hotspots, heating and, ultimately, blackouts.

In Quebec in 1989 a large transformer was blown up thanks to GICs. The province of Quebec lost power for nine hours and it caused millions of pounds' worth of damage, so, ideally, what we want to do is prevent that from happening.

If we know an event is coming there are steps that electricity operators can take to mitigate the effects. New Zealand had a plan which it had agreed with academics and they approved it last May, re-routed power, took things off maintenance, weathered the storm, and then claimed that the plan worked.

Direct GICs are rare and the equipment to measure them is not present in every transformer; they are expensive to install and there is no guarantee that future events would be predictable. We don't measure GICs but we do measure instead the rate of change of the magnetic field. This is the parameter that drives the GICs. To calculate this effect you need to account for the sub-surface conductivity to give you an estimate of the geoelectric field, and knowing the power-network configuration, combining the two allows you to calculate the GICs.

In our studies we concentrate on the geomagnetic field itself, and mostly use the rate of change of the magnetic field. If it gets to 50 nanotesla per minute you are going to get significant GICs. Our study is funded by NERC so we work on impacts on the UK. We have three permanent geomagnetic observatories in the UK from which we have data going back 40 or 50 years. The first is Hartland in Devon which has a current geomagnetic latitude of 48 degrees, with another at Eskdalemuir (55 degrees), and a third at Lerwick (58 degrees).

One process that causes GICs is called a sudden commencement (SC) which, in turn, is caused by the impact of a solar-wind shock in the magnetosphere. It causes a rapid contraction of the geomagnetic field. We have about one to one-and-a-half hours of warning which makes it forecastable. On the ground the magnetic field jumps by several hundred nanotesla in one to two minutes and the effect varies with latitude.

How much does this matter for us? Not much as it turns out. It depends on magnetic latitude and is more important at lower geomagnetic latitudes. However, in the days after an SC we find that we are accounting for 80% of our low-latitude station variability here. Less than 10% of larger magnetic variability is connected with SC events but more than 90% in that three-day interval after. There are two sub-types of SCs — sudden impulse (SI) and storm sudden commencement (SSC). If we see the interplanetary shock rapidly compress the magnetic field and nothing happens then that is classified as an SI. If we see a geomagnetic storm kick off after that, we call it an SSC. We find that the only events we need to care about are those which trigger a geomagnetic storm, so only those SSCs are significant and if we have a three-day warning then we will see the greatest threat to the power networks that we are going to see.

About 30 minutes before the shock impact we see the structure in the solar wind — the solar-wind density increases tremendously, the velocity with which it approaches us drops a lot, the magnetic-field strength increases, and there is a small rotation of the magnetic field so that it points northward. We extract features from the solar-wind parameters and use a series of models invoking linear, non-linear, and ensemble methods to see whether a storm or SSC is going to happen. We restrict the number of parameters we use — most important is the range in B, the magnetic-field strength, followed by the range in particle density and the range in solar-wind velocity. We calculate models using increasing numbers of parameters and we are interested in getting the model reliability — the parameters tell us if something will happen, or won't happen, but how much can we trust it? The second metric is called the skill and how good the model is from picking out events from non-events. We find that we can get all the information from the solar wind using three parameters.

We also ask is there going to be a geomagnetic storm? Again we rank the parameters by importance. This time the most important parameter is B_z, but this time we need seven parameters to describe the shock. What drives the geomagnetic storm and what follows the shock? The four statistical models we use give a different answer which is purely a consequence of the model. A more direct approach is to ask the model what we might see in the future. Instead of predicting the exact shape of a model event, they will be asked whether it will be above a certain level, *i.e.*, a threshold. Three models of increasing complexity were tested and placed into a neural-network model. From the metrics here we can see that the more complex models don't necessarily see the best performance.

In conclusion we note that forecasting GICs is an important space-weather problem. Magnetospheric phenomenon-based forecasting can give an insight and broad warning interval. Direct forecasting can produce a few hours of warning but less insight. We find that simple models perform as well as complex ones but extrapolating the model is dangerous.

Professor Ofer Lahav. Thanks for an interesting talk. I guess those two categories you mention: feature extraction as the more direct one, the second is an example of deep learning. Is this correct? I'm an outsider in this field but in other applications deep learning really does better, even if it is more difficult to understand what it does, so I wonder if you could explain why the deep learning is not doing better, relative to feature extraction?

Dr. Smith. For context, we went back last year and we used methods like SHAP (machine learning) to try and explain why the models are doing what they're doing. It turns out they are replicating the physical understanding that we have, in some ways. I'm still not entirely sure about their extrapolations. However, that's a slightly different question. I think the problem that we're having with deep learning here is that the system is massive and complicated and we're driving it with a single point of input, and that single point of input may not even be that reliable, honestly. It's upstream of the Earth but the entire spatial structure within the solar wind can be very complicated and there's no guarantee that what we're seeing is representative of the phase front of the solar wind as it comes in, for example. So, the models are having to interpret what's going to happen in terms of magnetospheric activity and there's a whole bunch of processes that could happen and then map down to the Earth. It's a complex model that's not fully possible to solve, even with things like numerical physics models of this case. So it's asking for a lot for the model to interpret that, and I think that's the problem.

Professor Tim Horbury. Really nice talk. I'm just going to give a shameless plug, if that's OK? One of the issues in terms of predictability is if you're in the sheath upstream of a cloud, it can be very variable, but if you actually drop into the cloud itself, it's much more predictable. There are ways of diagnosing that, for example, with temperatures or charge states. So that's something that you might want to think about putting into those kind of models. Charge states aren't available at the moment, but *IMAP* is launching in September, and that will have real-time charge-state data, and at that point, we're going to have four spacecraft at L1 with real-time data at once. So, as you mentioned right at the end of that answer, multi-spacecraft here will make a big difference in being able to make accurate predictions in real time, I think. I'm not sure many people are really thinking about it yet, but I think we should really get on and start thinking about how we combine multi-spacecraft L1 data to make better real-time predictions.

Dr. Smith. I absolutely agree, and we've got an ISSI team we put in yesterday to look at things like that, so I hope to review that.

Professor Lester. I'm wondering what would have happened if he said that he could do a shameless plug and I said "No!"

Professor Horbury. I've got the microphone and I'd have done it anyway!

Dr. Smith. The last point on that is that these models are all very data intensive. So, we were training these models on 30 years' worth of data. The problem is that we wouldn't have the historical charge-state data that we may require for that, so it might require some clever thinking maybe to go back and create some sort of substitute pseudo-data that would fill the gap.

Professor Whaler. Thanks very much, Andy. I've probably missed something in your talk but are there any advantages or possibilities of data assimilation from, say, measurements of the magnetic field on the ground, for instance, that might help with the types of models that you've talked about today?

Dr. Smith. Certainly, they definitely do that with things like the ionospheric models that they have. They incorporate data assimilation and that boosts the model performance massively. If we were predicting the ground magnetic field itself, then certainly data assimilation would be an excellent thing. Certainly over the past five to ten years, we're starting to build in more of the kind of techniques that have come in from weather forecasting, for example, leaning heavily upon our friends, such as at Reading. I think that's where we're going in the future.

The President. I'm almost shocked by how different the skill-score metrics are in there and what they tell you and which is the right one to use depends on the application. Are you using ones that are guessed to be the best for our network operators? You see what I'm saying?

Dr. Smith. Yes. I'm not sure if there's consensus at this point as to what we should be using. There are some sort of typical ones that people go to and then there's also reasons why they are potentially not the best ones, because of things like unbalanced data sets and we should be using something else. Mike Liemohn has a fantastic paper on metrics and which we should be using, but that mainly focusses on regression rather than classification.

The President. So it is interesting about metrics; I suppose this is true of life, the trick is to choose the one that's right.

Professor Phil Charles. Just a quick one, because I forget whether it was you or Mike who had the table of the top-ten solar-wind events. I was interested to note that 1989 wasn't on there, yet we know it did a lot of damage on the ground, and also, if I remember correctly, it took out several satellites that were operating at the time? How does that correlate with this work?

The President. It's because there's no such thing as a typical storm. A space storm with not a few particles that destroy your satellites and things is rather different from a geomagnetic storm. They do go together but they're not one for one. So, the classic example was 1972, just before the Apollo mission. If it had happened while the Apollo astronauts were out of Earth's magnetic field, it would have killed them. On the ground, that was a very small geomagnetic storm, but it was a massive space storm in terms of that. The only thing it did on the ground was all the American mines laying around harbours in Vietnam exploded and nobody knew why, but there wasn't a major geographic event. It's very varied, so that's part of the problem.

Dr. Smith. There's an entire zoo of processes and what makes a big storm in some senses may not be in others.

The President. Exactly. Again, which metric you used to filter.

The President. We move on now to the Eddington Lecture — 'Reconstructing the history of the Milky Way galaxy using stars'.

Dr. Melissa Ness. (It is expected that a full summary of this talk will appear in a future issue of *Astronomy & Geophysics*.) [The speaker began by stating that astronomy of the Milky Way galaxy has entered a transformative era. The *Gaia* mission and an ensemble of ground-based spectroscopic surveys are delivering element abundances and velocities for millions of stars. These data provide both an opportunity to deepen our understanding of galaxy formation and to test limits of 'the limit of knowledge'.

The speaker continued by summarizing the endeavour of galactic archaeology, whereby we are working to reconstruct the history of the Milky Way galaxy using signatures in its stars. Dr. Ness laid out the data landscape in the 2025 era and highlighted some key results to date from *Gaia* that showcase the differentiating behaviour of different stellar populations in the Galaxy, and then moved on to a discussion of the development of data-driven modelling techniques to derive information from stellar spectra, placing different surveys on the same scale, and talked about the new avenues that data-driven approaches have opened up. This includes enabling spectroscopic ages as well as individual abundances to be inferred from low-resolution, low-signal-to-noise spectra. At the heart of the talk was an exploration into the behaviour of the multitude of element abundances from the large surveys; how they are correlated, what amplitude of intrinsic information they carry, and how much diversity stars show in their abundance patterns (Fig. 1).

There have been several surprises that have come out of the large stellar surveys and data-driven methodologies built to analyse them. We have learned that up to 1 in 100 stars in the disc are 'abundance *dopplegangers*' — chemically identical but unrelated — limiting the prospect of reconstructing the disc's star-cluster building blocks. Furthermore, for stars in the disc, most of the element

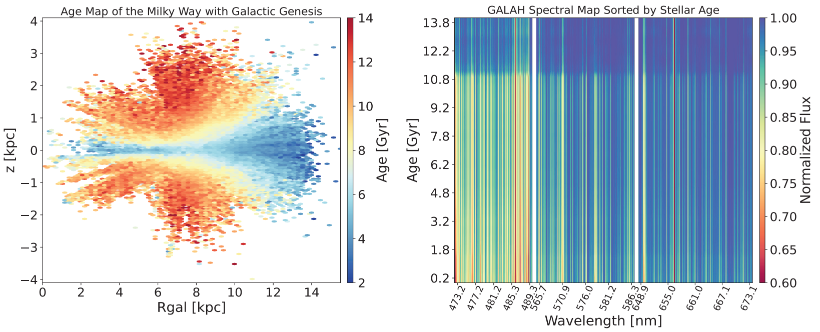

Fig. 1

The left figure shows the age map of the Milky Way, from the Galactic Genesis survey. The right image is a map of continuum-normalized *GALAH* spectra of red-giant stars, with wavelength along the x-axis and stars sorted in rows along the y-axis in order of ascending age; this summarizes the way that information is encoded in stars; the amplitude of features changes across wavelengths and it changes across stellar age. This is the information we access and interpret, to learn about the formation and evolution of the Milky Way galaxy and map the Galaxy across its spatial extent.

abundances measured for most of the stars can be predicted to a precision of better than 10 percent given only two key abundances. However, this is not the case for stars in the stellar halo. These findings frame how we can most effectively work with the data to turn photons into a quantified description of Galactic history and provide strong constraints on the star formation and mixing processes that have set the Galactic environment.

Dr. Ness concluded by summarizing the remarkable surprises and signatures we have found in the data, and highlighting prospects for chemical tagging in the halo and data-led nucleosynthesis in the disc with the upcoming data from new surveys.]

Professor Richard Ellis. What a beautiful lecture. Thank you so much. Your result on testing chemical tagging is very sad because it is such an elegant idea, but maybe we're being too picky. If you applied chemical tagging to the disc — I mean, with the test that you did: a cluster *versus* the field stars — are you honestly sure we wouldn't learn something just applying it to the field stars? Even if the map of geographical location is slightly blurred, surely there's some benefit and utility in this nice method?

Dr. Ness. I think the power of chemical tagging is not that it works, but it's the method, it's the approach, it's how we're learning so much from the data. So I absolutely think that chemical tagging is incredibly useful as a tool to learn about the data, and I think there is much to do in the abundance space by looking. With the abundances we learn we've basically traced stars back to their birth radius, not to the individual clusters. So the abundances are very discriminative for that. So I'm very interested in taking this to its limits. It's more the theory of chemical tagging — it doesn't matter whether it's successful or not. It helps us enormously.

Professor Claudia Maraston. Thank you very much for this lecture. I'm going to try to export your wisdom to external galaxies, for which we derive metallicities up to very high red-shift from low-resolution spectra. And people, I mean including us, we have done this to attempt to reconcile very-high-resolution abundances to low-resolution abundances. You had a couple of points, but I failed to get the conclusion. When you compare low resolution, like *LAMOST*, with high resolution, what was the result for most elements and also comparing analysis from optical to near infrared, where you notice offsets, which one is better?

Dr. Ness. I won't get into which is better in the sense that we tend to adopt the high-resolution values because they're more precise and I suppose better and they will have different surveys, your *LAMOST* and your *APOGEE* and your *GALAH* have different scales, but if we use a data-driven model and use the styles in common we put them all on the same scale: we have to adopt. So you have to make a decision about which one you want and that's the big question. Which is the more accurate. I think there are tricks we can do; because we have a generative model it should fail to be able to generate the spectra with the wrong labels. I think you can test the fidelity of the labels, which set of labels of which survey is more correct by seeing where you fail to generate a good model of the data. We can talk more about that idea. I'd like to try that. I think it would be fun.

A Fellow. Thanks for the nice talk. I am interested that there are rare events in your list of enrichment sources and less-rare events. So do you see a larger star-to-star scatter in elements that are expected to come from rare events?

Dr. Ness. Good question. So you mean the neutron-capture elements, like some things from r-process?

The Fellow. Yes, something from the neutron-star merger should be rare and I wouldn't expect the whole galaxy to be enriched in the same way by neutron-star mergers.

Dr. Ness. I think there are two points. Yes, all of the neutron-capture elements have higher scatter, but we are predicting on non-neutron-capture elements. I think in the residuals, the fact they're all so correlated, implies that they're all generated from the same underlying 'other source'. I think that we can do that test, but to do that we need really to look at a range of metallicities and look at how many stars are failing our prediction — about 1% fail. In particular, we see things fail in barium and yttrium because that's produced, I think, in stellar-companion mass transfer. I think that we can test that. We just haven't yet.

The President. We should thank Melissa again for a fabulous talk [applause]. I have some questions for you, but I will ask you later today. The next meeting will be 11th April in the Royal Irish Academy, Dublin.

A SIMPLE, STABLE, RAPIDLY-CONVERGING, AND EXTREMELY ACCURATE ITERATIVE SOLUTION FOR KEPLER'S EQUATION

By B. Cameron Reed

Department of Physics, Alma College, Michigan

A very straightforward scheme for iteratively solving Kepler's equation is described. The iteration method is the familiar Newton-Raphson technique, augmented with an initial solution estimate based on modelling the sine function as a downward-opening parabola. For even extreme eccentricities, micro-arcsecond accuracy can be achieved with about a dozen iterations. No use is made of 'canned' calculator or spreadsheet root-finding algorithms.

It is well known that it is impossible to obtain a closed-form solution of Kepler's equation for the eccentric-anomaly position angle E of an orbiter as a function of time t. This problem has accumulated four centuries of analysis, and a student or non-specialist researcher who explores the on-line, textbook, and journal literature on it will soon come across intimidating-looking series expansions with coefficients involving Bessel functions, dire warnings concerning convergence instabilities, and a plethora of strategies for generating an initial estimate for E.[1,2,3]

Solving Kepler's equation is now trivial given the root-finding routines available in calculators and spreadsheets. As an instructional strategy, however, defaulting to 'black box' solutions whose inner workings are opaque has obvious disadvantages. The purpose of this note is to describe a very simple,