Driving Net Zero

is agritech ready to capture carbon?

cambridgeconsultants
Part of Capgemini Invent

Contents

1 Executive summary

The regulatory and subsidy systems underpinning farming are undergoing radical change. Furthermore, issues such as soil health, net zero and farm nature value are increasingly in the public eye. A key change to be expected is a transition to payments based on *outcomes* rather than a traditional agricultural balance sheet. An important outcome could be carbon sequestration accomplished by changes to farming practices.

In the case of carbon sequestration, payments will need to be based on the amount of carbon taken permanently into the soil, which will require cheap, accurate measurements. Currently this is not possible without manual sampling and centralised lab analysis. This report shows how modern techniques proven in other industries, coupled with modern data science techniques, can be used to measure overall carbon uptake in an economically viable way.

The benefit of this would be to democratise these measurements, opening them up to a sufficient number of farms to make a real difference. Agritech industries would then be able to lead these new markets, selling products and services which fit these new financial regimes.

2 Agriculture and the carbon market

Agriculture is the only industry which can easily take carbon from the atmosphere in quantity without the need to develop huge new infrastructure. Plants do this naturally – taking CO_2 in and storing the carbon in their cells. If plant material can make its way into the soil, then carbon can be removed from the atmosphere, both improving the soil and contributing to carbon reduction targets. If this carbon uptake in the soil can be measured in a reliable, acceptable method then the farmer can participate in a carbon market and in turn receive extra income. This income can pay for improved land management practices, which then create other 'public goods' which will fit future subsidy regimes.

Carbon sequestration covers a variety of practices such as biochar (where carbon rich products are used to improve the soil) and enhanced weathering, where a quarry product such as basalt is added to the soil. Of greater interest is the use of modern land and soil management practices to increase the natural ability of plant's lifecycles to absorb carbon from the atmosphere and retain it in the ground.

Receipt of payments for removing carbon at least semi permanently will require detailed measurement, auditing and a convincing argument that these benefits would not have happened anyway. These measurements naturally fit with another change in agriculture – the move to outcome based financial models.

"If this carbon uptake in the soil can be measured in a reliable, acceptable method then the farmer can participate in a carbon market and in turn receive extra income."

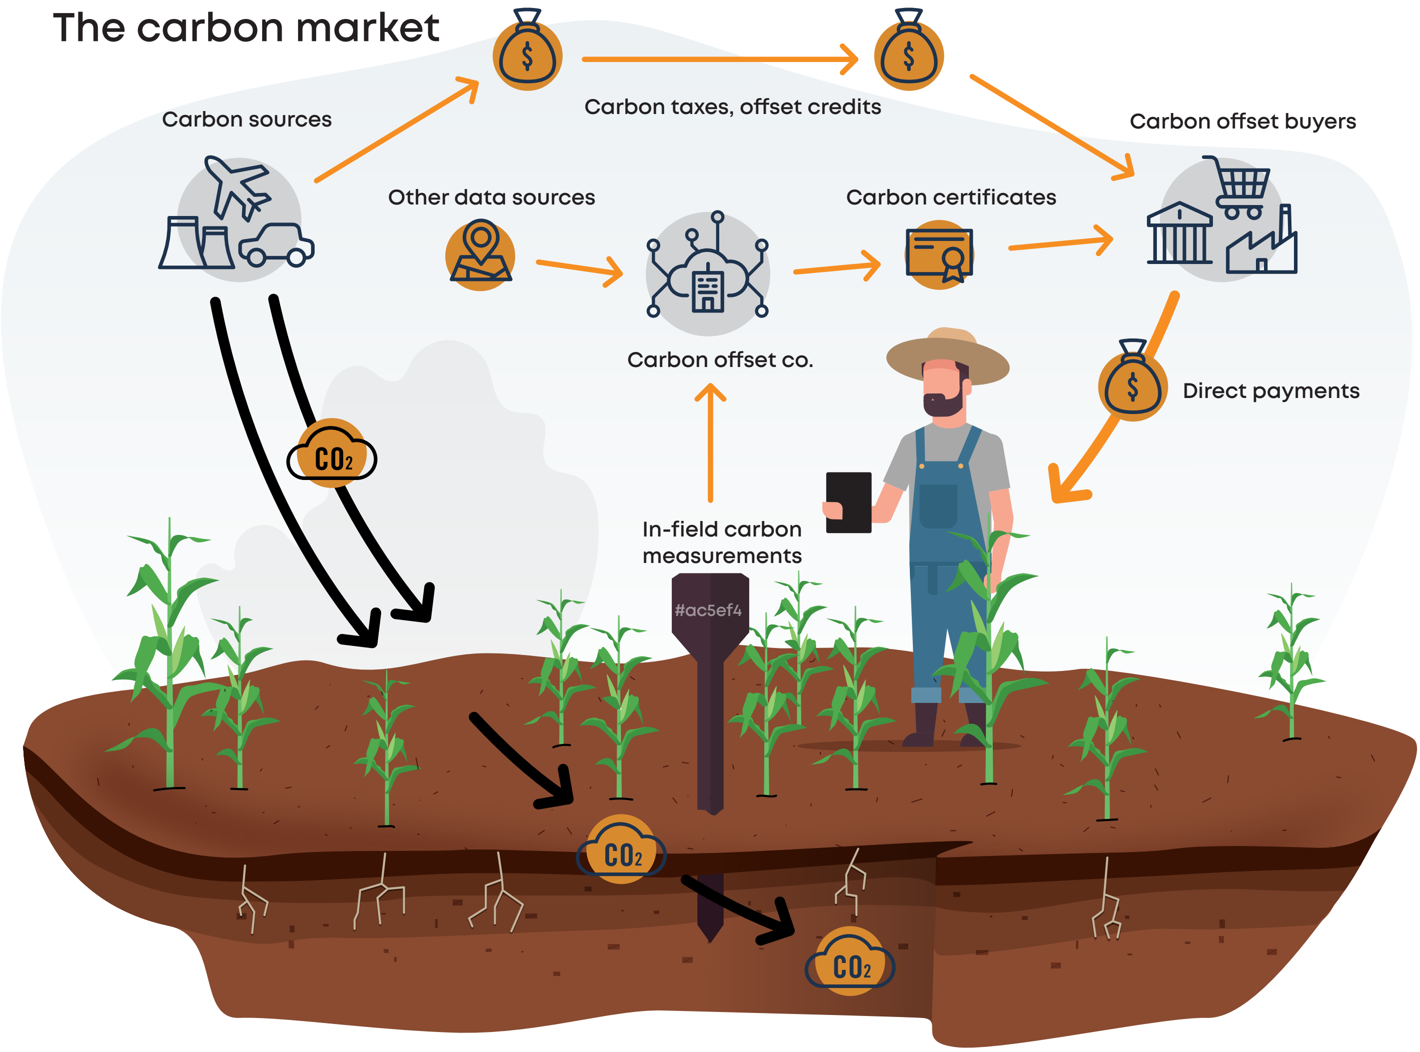

Figure 1: The carbon market

Outcome based agriculture

- Farming is the only industry which is ready to become carbon negative – taking CO2 out of the atmosphere and burying it for the long term
- This is done with updated land management practices – but needs verification and a chain of trust
- Agriculture is moving this way anyway – ag co's are selling *outcomes* rather than product, so linking in CO2 sequestration is a natural next step

Chain of trust

- Needs to be believable – measurements are backed by proven lab tests
- In-field tests have to sample the whole farm in a representative manner with an economic number of samples
- Only changes slowly, so continuous monitoring required – but this links to many other ag initiatives

Benefits to farmer:

- Direct income from offset payments, public investment in updated practices
- Healthy soil catches more carbon, carbon makes soil healthier (so reduced inputs needed)
- Increased biodiversity which future subsidy schemes will reward

3 Outcome based farming

The drive towards digital agronomy and connected machines has led to a wealth of data being available from on-machine sensors, remote sensing and in-field measurements. Although this data is often siloed between different proprietary interests, bold initiatives are arising which aim to connect these different data sources. The interest is not academic – sharing data up and down the chain can improve outcomes for everyone from seed suppliers to flour millers. Indeed, this is one of the reasons why much of the interest is driven from large consolidated food producers and retailers. Importantly, it also allows financial risks to be shared – the costs for some farm inputs are dependent on a particular outcome being achieved.

To achieve this outcome, and price carbon correctly, field conditions need to be monitored throughout the growth cycle. This isn't just about ensuring compliance with the product requirements, but to be able to give help and advice. This need is driving a range of sensor developments, for instance smart insect traps, climate sensors and particularly in-field soil sensors. It's also driving a need to connect up all these sensors and machines, using Internet of Things (IoT) techniques, borne over 5G and specialised radio systems suitable for the remote agricultural environment such as LoRa and NBIoT.

Remote sensing also has a part to play – for instance high resolution satellite or drone images, particularly when interpretation is aided by AI techniques. However, carbon is stored deep (30cm) in the soil, and it can't be measured directly remotely. It could possibly be measured indirectly however, by modelling the expected carbon uptake by plants at particular growth stages and counting them remotely. Although this model would be complex, it could form a low accuracy but high sample size across a wide area (see Figure 2).

"The interest is not academic – sharing data up and down the chain can improve outcomes for everyone from seed suppliers to flour millers."

4 Measuring soil carbon

Soil carbon is a complex collection of plant matter and the remains of complex biological processes. Measurement of 'carbon' has historically been performed at a lab scale by pyrolysis – burning a dry soil sample and analysing the gases given off. This process is slow, labour intensive and requires significant skill in the sample preparation and analysis. If the industry is to scale up these measurements to a wide land area and really play its part, then it needs a quick and cheap test which:

- Is 'accurate' in that it has an excellent correlation with historical and lab measurements
- Accounts for bulk density (which is currently a common and large source of inaccuracy)
- Can sample at depth, and properly account for the layering and depth profile of soil
- Is economically effective: i.e., an unskilled person can get an effective number of samples in a work day

These measurements, made at ground or root level, connect into a wider data aggregation system which makes the overall estimate, as shown in Figure 2.

There are two steps to achieving this: making the measurement itself; and interpreting the reading back to give the quantity of carbon across large areas of agricultural land. As such several pieces of information are required: stratification of the soil based on topology and underlying geology to inform sampling; measurements of the soil carbon at specific locations; and measurement of the bulk density of the soil to inform the extrapolation over the landscape.

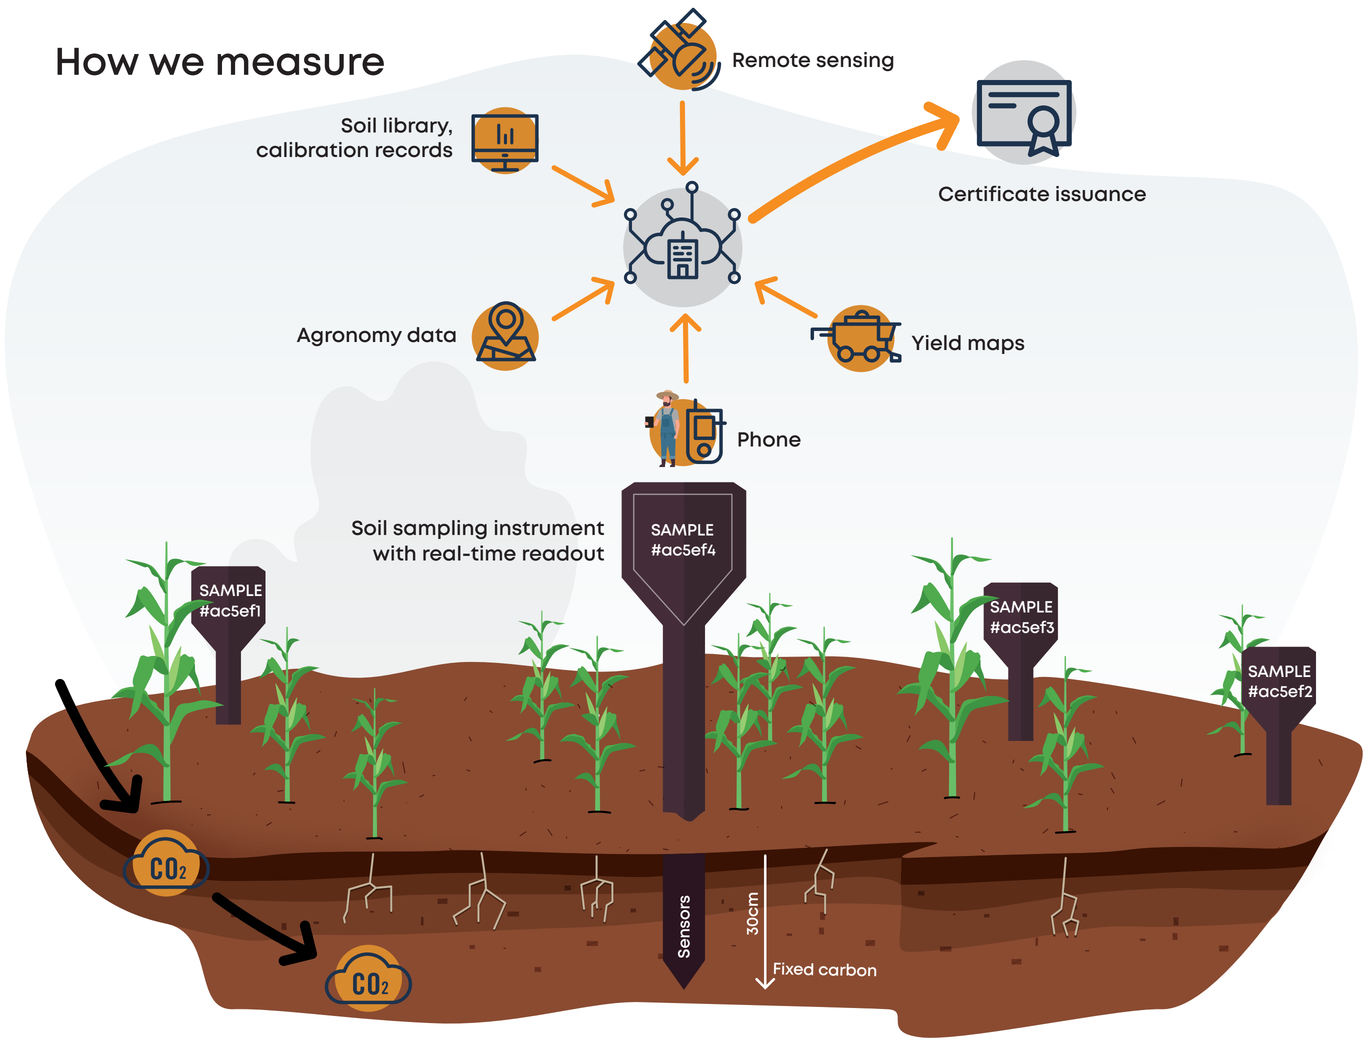

Figure 2: How we measure

Why test in field?

- Important carbon capture processes happen below ground – surface measurements can't tell you everything
- Land use is complex – every few metres could be different – satellite or drone is a quick way to survey this area
- But we can *calibrate* the images for carbon if we have enough in-field measurements
- Key is use data science to merge many streams – from once per year lab tests to detailed yield maps deriving a believable figure

Certificate issuance

- Auditable – did these things really happen?
- Creditable – were these measurements carried out properly and interpreted correctly?
- Backed by science? Good correlation with gold standard lab tests

On farm devices

- Guides through where to sample based on remote sensing data, history and where changes are expected
- Gives instant feedback that the device is working correctly
- Optimised for low / no data usage outdoors

4.1 The sensor

The requirements on this sensor – simple to use, cheap and effective – are similar to those on a medical diagnostic or screening test. The medical industry has made great strides in taking lab based tests (slow, resource hungry) to point of care where near instant, clinically accurate results can be obtained. These tests are built typically on novel chemistry and cutting edge optical sensing. A good example of this are lateral flow tests, which are able to prepare the sample, carry out the immunochemistry and read out all in a low cost disposable device.

In medical systems, the chemistry is well understood, and typically highly consistent between individuals (we all have the same biology). However, in the case of soils, we cannot expect the same level of consistency. Not only are there different land uses (e.g. arable, pasture, horticulture) but a range of soil types, underlying geologies and climates.

There are many models for the deployment of sensors for soil carbon sensing, each of these has its own advantages and trade-offs. Example models include:

- A sensor located in a laboratory, to which soil samples are shipped. This is largely the existing paradigm and has serious limitations on sample throughput and on the costs associated with shipping large numbers of samples to a central location. This may be significant due to the huge areas of land that have to be sampled to get landscape-scale measurements of soil carbon, however this model does offer the highest carbon measurement accuracy per sample
- A small mobile sensor can be deployed to the measurement location to collect a large number of measurements. This requires a durable and portable sensor with minimal sample handling, but may strike a balance between accuracy and the higher number of samples given by a centralised model

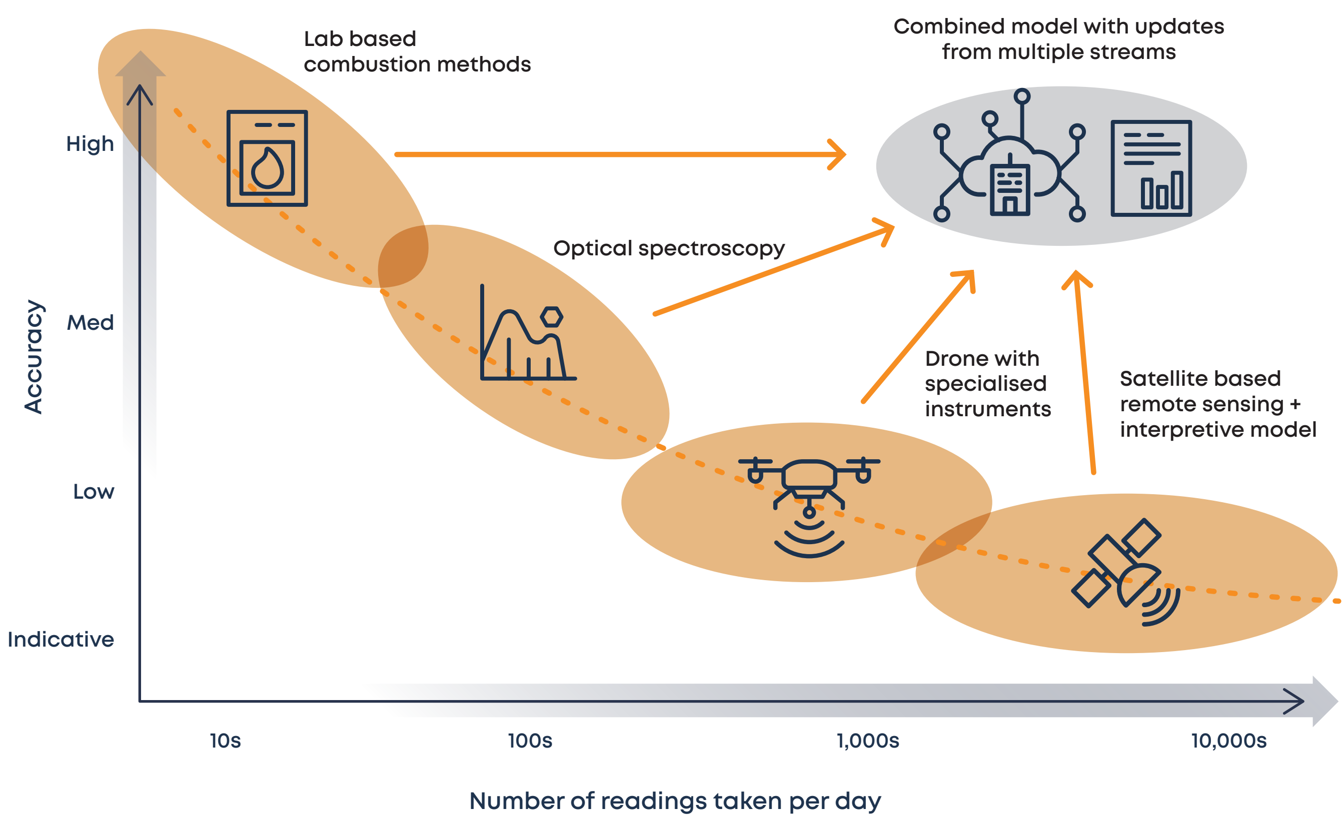

Figure 3: Different sensing techniques give access to different measurement accuracies, but the demands of sample handling affects the number of samples and amount of coverage that can be achieved

- Remote sensing from an aerial or satellite-based device offers the opportunity to collect huge numbers of samples over very large areas and can be highly automated to minimise operator time. However remote sensing is limited by the available pixel size, interpretation of vegetation cover over the ground and lack of penetration into the soil

- A sensor that is implanted in the soil and takes continuous measurements at a point location. While this system offers an effective method of tracking changes over time, the sensor has to be low cost (due to the large number required) while also being extremely durable in a hostile environment, having a robust communication equipment and a multi-year lifetime. It is likely that the design requirements of such a system are too demanding for such a system to be realised

All the above models require a robust chain of trust regarding the location of the sample measurement and the quantities of carbon recorded. In practice, it may be necessary to combine data from a number of different sensors of different accuracies to form a coherent model for soil carbon as shown in Figure 3.

Thermal combustion methods, such as loss-on-ignition, elemental analysers or pyrolysis, are considered the gold standard for soil carbon analysis. They also have the benefit that all past measurements have been made by them, making them the comparable for historical modelling. These techniques ignite gram-quantity samples of dried soil at high temperatures and analyse the mass changes and the composition of the resulting plume. Depending on the exact technique used, the quantities of organic, elemental carbon and possibly artificially added carbonaceous materials can be determined as a fraction of the total dry mass. Accurate bulk density measurements are critical to extrapolate from the combustion results to total soil carbon in a given area.

Combustion methods demand that a soil sample, such as a core, be collected then: segmented; dried; homogenised and only afterwards subjected to combustion. Therefore, there is a significant sample processing overhead to using thermal combustion techniques, most of which is currently done by skilled technicians. If a combustion method is to be a suitable technique for landscape-scale analysis, the throughput of sample processing needs to be dramatically increased over that currently available. This will require novel automation techniques to reduce sample preparation time and cost overheads. Without such as system, combustion methods will be prohibitively expensive to perform on sufficient scale to monitor entire landscapes.

The high processing requirement for combustion methods, even with suitable automation, means that they are largely only suitable for use in a centralised laboratory with core samples shipped from the site of interest. Unfortunately, the costs and labour requirements for digging cores and sending them to a lab do not scale as well as on site measurements.

For the measurement, optical sensing also fits the bill, as it's believable that it can be fast, robust and able to work without consumables or demanding maintenance. The ease of operation of spectroscopic techniques means that they can be readily produced in a form that can be taken to the site of measurement and collect many data points. This type of measurement, optical spectroscopy, has advanced in many industries. Scanning from the infra-red into the visible spectrum can be accomplished with off-the-shelf devices; the challenge is more to fit them to a form that can work reliably and repeatably in the soil, and demonstrating that consistent results are obtained.

How to interpret the spectra produced is less obvious. The problem is that carbon exists in many forms which will all present different spectra. The background from soil will also vary, to complicate the picture further there will be variable moisture levels and soil densities. Typically, industry solves this with a range of machine learning and AI techniques. The idealised workflow is 'simple': take a range of samples, measure the carbon levels in the lab to give ground truth and then train the AI to predict the carbon level from the spectra. However, there are many obstacles to overcome in this process.

Regardless of the soil carbon sensing techniques, extrapolating from a given sample to a landscape scale requires an estimation of the bulk density of the soil. Therefore, bulk soil density measurements are critical to accurate estimation of soil carbon across a large area. A number of methods to estimate the bulk density exist, but directly measuring the density of a soil core of known volume is often the most reliable method. However, it requires a large sample volume and significant sample preparation. And due to the large variability in soil composition and the multiplicative nature of the calculations, errors in the bulk density can lead to wildly misleading estimations of total carbon stored.

"Bulk soil density measurements are critical to accurate estimation of soil carbon across a large area."

4.2 Sampling

Sampling size is dependent on the accuracy we want to achieve, measurement device errors (precision) as well as real variation of the carbon in the field.

Soil carbon varies with time (i.e. growing season) as well as in field variation (as shown in Figure 4) and with depth. All these would contribute to the errors per measurement. Knowing the largest sources of error will tell us the sampling size we need to achieve the required accuracy. This is key: the overall error budget for the farm scale measurement – even a highly accurate technique such as lab pyrolysis can be let down by insufficient sampling.

Figure 4: Yield map showing spatial variability of a crop

If there were no errors in measurement and the carbon were the same throughout the field one measurement would suffice. But:

- How do we know where to sample and how many samples we need?
- The variability of the in-field variation will determine the minimum sample spacing to 'catch' all the detail likely to be present
- The more we know about the model the fewer samples we need and the greater the emphasis on measurement accuracy. If we know less about the model the emphasis will be on more sampling
- The more samples we take the greater the accuracy, with higher sampling rate driving cost and with diminishing returns
- Therefore, the decision on the sensor type and the solution formation is strongly linked to the model we assume represents the underlying carbon distribution and evolution

Sampling can also change over time. If many accurate measurements are needed at the start, key measurement areas can be identified by modelling thus reducing the number of measurements or moving to less accurate more cost effective solutions. There are many ways in which modelling the field variation or historical sampling can lead to optimising either number of samples or reaching the best sampling positions.

4.3 Interpreting the readings

It's a given that simply trying a range of machine learning algorithms on a properly organised dataset will yield some kind of result, giving a reasonable estimate of the lab carbon measurement from the in-field measurement. But, the problem with these predictive tools is *generalisation.* For instance, if we train an AI to recognise animals – what happens if we show it an animal it's never seen before? Or merely in circumstances that weren't expected (for example if the picture is upside down).

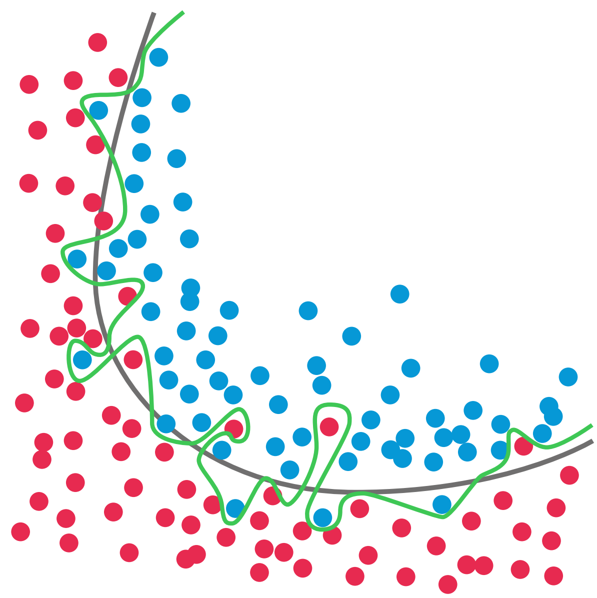

Figure 5: Example of 'overfitting' where the green line excessively tries to take in every point. The underlying model is actually the black line

This is the point where the *science* part of data science becomes important. For instance – if the algorithm was trained using pasture in the UK, will it work as well in the US? Or in Australia where conditions are arid?

The only way to achieve success will be to slowly stretch the breakthroughs we have today to new circumstances, carefully scaling up and learning as we go. This will mean taking in a wide range of data sources from the beginning – historical soil libraries, lab measurements and new in-field data.

5 The last leg – where to dig?

The previous section explains that since you can't sample everything, *where* you sample suddenly becomes important. For instance, if you need a random selection of people it's not enough to take numbers from a telephone directory and ring them. That will bias the sample towards those who have a phone and are at home to answer it. Similarly, where we stick the probe on the farm will be important: some areas may have dramatically high or low carbon levels, and bias the overall reading we arrive at.

The best way to decide where to sample is to have an underlying model, which contains what we know about where the carbon is, and which directions it's likely to move in. That way we can sample the areas where something is likely to have changed.

5.1 A digital twin for farmland?

A digital twin is where a detailed model of an industrial system is built up, with a wealth of sensor and operational data. That way its status is always known exactly, and 'what if' scenarios can easily be played out. The difference is that an artificial system like this tends to have a detailed and accurate model behind it such as the simulations carried out to design it correctly. Soil is a lot more complicated – for instance it depends on multiple, interlocking biological systems such as bacteria, nematodes and earthworms.

However, we don't expect to understand everything – just 'enough'. In fact, in some areas, precision agriculture has brought us to 'enough' already. Nitrogen, irrigation and pest treatments can be coordinated by in field and remote sensing images – and a similar approach should be viable with carbon.

In fact, some jurisdictions are already estimating carbon in a 'hands off' manner – purely by auditing what has been added to the field and how it's been managed. This has been backed by experiments and trials, making it believable, but depends on the collection of accurate and honest data. This data needs to be validated, archived and fitted to an underlying model by a central authority, so as to observe changes happening over time. Several of these bodies now exist due to government programmes or private investment.

So, the overall solution is to build two models – one at the level of the soil (covering several layers), and another which covers a few hectares. By keeping pace with changes that are made with a range of sensor inputs (remote sensing and agronomy data) this model can be kept up to date, predicting the level of carbon which the soil is holding. Over time, these models will drift and become less certain, which is when in-field measurements are required to 're-anchor' them back to validated ground truth measurements. This can also spot any attempted gaming of the metrics.

The model of the soil will enable to convert the sensor reading to a mass of carbon, including bulk density, corrections for moisture and other variations.

The model field will feed into the sampling to advise where to measure and enable a more accurate kg Carbon / hectare estimate.

The two leading measuring methods are pyrolysis and spectroscopy. Each have merits and the choice of measuring method will come down to geography and history as well as accuracy. Pyrolysis may be favoured, being the more accurate per measurement currently can not be measured in field field with fast turn around. Another reason for favouring Pyrolysis is historical, all previous data collected on soil carbon has been measured in this method.

Fields and soil types with large historical datasets will benefit from measuring by the same methods. Since all historical data has been taken with pyrolysis this biases the solution towards it.

However, some geographies or on continent level sizes will have difficulty to maintain pyrolysis based measurements. If there is little historical data and all modelling data needed to be taken quickly, a fast/cheap measurement regime is needed. For these two cases spectroscopy or a combination of pyrolysis and spectroscopy solutions are more suited. Although these measurements use quite different principles, they can be linked and interpreted together scientifically, particularly by using the modelling steps described above.

"Nitrogen, irrigation and pest treatments can be coordinated by in field and remote sensing images – and a similar approach should be viable with carbon."

6 Conclusions

Of primary importance to the carbon market is the 'chain of trust'. If I buy an offset credit, do I know that this carbon has really been buried, and can I be sure that this wouldn't have happened anyway? Most of the technology to bring this chain into being is already in existence – but it will be a combination of science and engineering which will link these elements in a way which can be trusted. Although connectivity will be important (bringing techniques such as IoT and 5G to the countryside), it's about connecting the *trust* between these elements and different measurement types with carefully validated measurement schemes and instruments.

The time is ripe as subsidy frameworks and incentives are changing worldwide, bringing new metrics such as soil health and biodiversity to farmer's balance sheets. These have similar measurement challenges and can benefit from a similar approach. These new methods to monitor land at both large and small scale will benefit us all, but only if implemented in a fair and transparent manner that brings benefits to the farmer and the entity conducting the 'off-setting', but without bamboozling both. In short, we need rapid, accurate, decentralised testing, overseen by the appropriate accreditation and regulatory bodies. The building blocks are there; now is the time to fit them together.

About the author

Simon Jordan, Head of Industrial Sensing,
Cambridge Consultants

Simon's role is to align commercial and technical requirements, and to develop ground-breaking sensor solutions for industry. He has worked across sectors such as navigation, oil and gas and agriculture, and is now developing applications for quantum sensors.

Contributors

Noga Sella, Senior Engineer
Alistair MacNair, Senior Chemist

About Cambridge Consultants

Cambridge Consultants has an exceptional combination of people, processes, facilities and track record. Brought together, this enables innovative product and services development and insightful technology consulting. We work with companies globally to help them manage the business impact of the changing technology landscape.

We're not content to deliver business strategy based on target specifications, published reports or hype. We pride ourselves on creating real value for clients by combining commercial insight with engineering rigor. We work with some of the world's largest blue-chip companies as well as with innovative start-ups that want to change the status quo fast.

With a team of around 800 staff in Cambridge (UK), Boston, San Francisco and Seattle (USA), Singapore and Tokyo, we have all the inhouse skills needed to help you – from creating innovative concepts right the way through to taking your product into manufacturing.

For further information or to discuss your requirements, please contact:

Niall Mottram, Head of Industrial & Energy
niall.mottram@cambridgeconsultants.com

cambridge consultants
Part of Capgemini Invent

UK — USA — SINGAPORE — JAPAN

www.cambridgeconsultants.com

Cambridge Consultants is part of Capgemini Invent, the innovation, consulting and transformation brand of the Capgemini Group. www.capgemini.com

Lit-R-015 v0.2 '2022_A2A_Driving Zero_Report'