


Calibrating a capacitive moisture/water content sensor for hydroponics

As I've mentioned multiple times in my blog, moisture sensing is one of the most important measurements in a hydroponic crop that uses a significant amount of media. It allows you to properly time irrigations and avoid over or under watering your plants. Capacitive sensors are the lowest cost initial approach to adequate moisture monitoring of your media. In this post, we are going to learn how to use an Arduino with a gravity shield and a low-cost capacitive moisture sensor in order to accurately monitor the water saturation of our media.



A capacitive sensor plugged into an LCD shield and an Arduino UNO. The measurements are very easy to track with this setup.

An analogue capacitive moisture sensor [like this](#) one does not expose any metallic parts to the media and can be used for the monitoring of moisture content. This sensor is powered by 3.5-5V and gives you a voltage signal that is proportional to the dielectric constant of the media it is in. As the dielectric constant of media changes with the presence of water, so does the signal of the sensor. However, no specific voltage corresponds to a specific water content measurement by definition, so we need to calibrate the sensor in order to interpret the voltage values we read from it. In order to get readings from this sensor, I use an Arduino UNO, the above-linked capacitive sensor, and an [LCD shield](#) from dfrobot that I can use to easily get the device readings. The arduino code used for this device is also shared below.

```
#include <U8glib.h>

#define XCOL_SET 0
#define XCOL_SET_UNITS 50

U8GLIB_NHD_C12864 u8g(13, 11, 10, 9, 8);
float capacitance;

void draw() {

    u8g.setFont(u8g_font_helvB10);
    u8g.drawStr(0,21,"M_SENSOR:");
    u8g.setPrintPos(XCOL_SET,51);
    u8g.print(capacitance);
    u8g.drawStr( XCOL_SET_UNITS,51,"mV" );
}

void setup() {
    pinMode(LED_BUILTIN, OUTPUT);
    pinMode(4, OUTPUT);
    Serial.begin(115200);
    analogReference(DEFAULT);
    Serial.println("MOISTURE");

    u8g.setContrast(0);
```

```
    u8g.setRot180();  
}  
  
void loop() {  
  
    draw();  
    capacitance = analogRead(1);  
    Serial.println(capacitance);  
  
    u8g.firstPage();  
    do {  
        draw();  
    }  
    while( u8g.nextPage() );  
  
    delay(1000);  
}
```

The calibration of a moisture sensor usually requires the creation of what soil scientists call a “water retention” curve, which is a curve where you plot the sensor’s signal as a function of a fixed volume or weight of water added to the media. However, this approach involves the use of many different containers and the addition of water to the media followed by oven drying steps in order to accurately determine how much moisture was actually in the media for every measurement carried out.


In order to do this procedure in an easier manner, losing as little accuracy as possible, I have created a calibration procedure that makes use of natural drying and only requires one single oven drying step. The procedure is as follows:

1. Heat the media that will be used for an experiment in an oven at 110°C (this drives out all water).
2. Wait for the media to cool to room temperature.
3. Fill the container that will be used for the test (this can plastic but has to have holes at the bottom). Put the sensor in the media, make sure the sensor is driven into the media until the top line is reached.

4. Take a reading, this is the “completely dry” media reading.
5. Disconnect the sensor from the Arduino.
6. Weigh the container+media+sensor. This will be the “dry weight”.
7. Take the sensor out.
8. Add water to the media until there is ample runoff coming out of the bottom.
9. Wait until no more runoff comes off.
10. Put the sensor in the media, making sure you drive it in until the top line is reached. The sensor won't be taken out for the remainder of the experiment.
11. Connect the sensor and take a reading. Take note of this value.
12. Disconnect the sensor.
13. Weigh the container+media+sensor. Take note of this value. The first reading you take is the “max saturation” weight.
14. Repeat steps 11 to 13 every 1-6 hours (time is not very important as long as you gather enough data points) until you reach close to the dry weight. This can take several days depending on the media used. The more points you measure, the more accurate your curve will be.

After carrying out this procedure, you can create a curve like the one shown below. You can use the difference between each weight and the dry weight divided by the difference between the weight at max saturation and the dry weight in order to calculate the water saturation percentage. As you can see, the curves for these sensors are fairly linear in the 40-100% moisture range for the media I tested, while below 40% the regime changes and the measurement increases exponentially until it reached the “dry weight” sensor value. The entire curve can be described with a power-law equation. This behavior will not be the same for all media tested, reason why it's very important for you to calibrate the sensor for the

specific media you want to use.



Calibration curve for a capacitive moisture sensor. In this case, the media was a mixture of 50% river sand and 50% rice husks.



Once you have your sensor calibrated you can know what a measurement in mV implies in terms of water saturation for a given media type. This allows you to time your irrigations at a given water saturation % effectively. Which water saturation percentage might be better, depends on the properties of your media and how the water potential changes as a function of the water saturation. However, this allows you to experiment with irrigations at different water saturation % values and figure out exactly where you need to water so that your plants are not under or overwatered.

It is also worth noting that the above sensor is probably not rugged enough for use in a hydroponic setup without a bit more hardening. In order to use these sensors in practical applications, you should apply conformal coating to the electronics at the top of the sensor and then use shrink

tubing in order to fully protect these electronics from the elements.

Properly positioning temperature and humidity sensors in a hydroponic growing environment


Temperature and humidity are two key variables you need to measure. They are important because they determine how your plants will transport water, and transpiration controls a lot of processes, with nutrient transport being one of the most important ones. However, the value of these variables in a growing environment – being that a greenhouse or a grow room – can change substantially depending on where they are measured. It is therefore critical to know where to place sensors and how to interpret their readings based on their location. In this post, we are going to discuss where it is best to measure these variables and what consequences it could have if these values are not measured properly.



Sensor placement relative to a control source (AC, humidifier/dehumidifier) for one or two sensors. Note that this setup assumes good circulation throughout the room, including middle of canopy.

Let's start with the worst possible case, you only have one set of sensors and you need to control your environment with it. *In this case, place your temperature and humidity sensors at canopy height, as far away as possible from both the AC ducts and the humidifier/dehumidifier, make sure the sensors are hanging in the air and not stuck to a wall or tubing.* Then, make sure you use a hot wire anemometer to verify that your air movement speed is at least 0.3m/s across the entire room. This setup ensures that the worst controlled part of the room is at the correct value and it also attempts to minimize the gradient created from the control sources to the sensor using a good amount of air circulation. It is not perfect though and significantly different "climate zones" will be created close to and far away from the climate control devices.

The above setup can be used effectively in small growing environments, but can be problematic as both the number of plants and the size of the growing environment increases. At this point, using a single set of sensors is not an option if adequate climate control is desired. In these cases, multiple sensors need to be placed to ensure that climate control is being done properly. When using multiple sensors, place the second sensor at the place with the lowest air circulation inside the room, at canopy height, which is usually in the middle of the room, then place subsequent sensors as far away from either this or the first sensor in sequence. When doing climate control, the system needs to ensure all of the sensors remain within a "safety band", making sure no sensor becomes too cold/hot, humid/dry, during control cycles.



Sensor placement for multiple control sources. Sensors are placed in order trying to always be as far away from sensors as possible but within the plant canopy.

When you implement a sensor system like this, you will realize pretty quickly that climate becomes very hard to control in a larger room when there is only one source of control (one AC, one humidifier, and one dehumidifier) because gradients become too big for effective control, so it takes too long for the AC

to be able to properly control the room while ensuring all sensors remain within proper boundaries. In this case, it becomes necessary to add multiple sources of control, so that the extent of gradients within the room can be minimized. This means adding multiple ducts for the output of an AC, multiple humidifier/dehumidifiers, etc.


In these cases, sources of control are placed outside of the plant canopy to avoid plants being exposed to the flows from these control sources (which would expose them to very cold/hot/dry/humid air). Then the sensors need to be placed within the plant canopy, starting from as far away as possible from all sources of control – usually the middle of the canopy – and then to the corners of the growing environment.

Note that the control algorithm needs to ensure all of these sensors are within the proper control band and not attempt to control the average reading of these sensors. If you try to use the average of sensors to control a room, you might be left with a room where two extremes are present and the control system believes everything is ok while these extremes are maintained. The median is a better way to control a room, but it only becomes useful when 5 or more sensors are used. If only 2 or 3 sensors are used, ensuring all of these sensors are within adequate bands is fundamental to ensure that the room will have a lower chance of having humidity/temperature microclimates that will be detrimental to plant growth.

Making the most out of your hydroponic setup's logged

sensor and control data


If you have a hydroponic crop with a data logging and automated control solution, you probably have a lot of sensor and control data recorded that could be useful to take your crop's results to the next level. In this post, I am going to talk about some things that you could be doing with these stored data. You will see how the usage of this data opens up many possibilities and that even implementing the most basic of these suggestions could lead to important improvements in your understanding of your crop and its results.



Use of different moving average to smooth out sensor readings. The lowest hanging fruit to take advantage of logged data is to be able to download the data and put it into a database structure that you can properly query and search. Most data logging solutions record the data in either very simple structures, like csv files, or non-relational databases – like MongoDB – which are rather limited and do not allow for the degree of versatility that a true relational database engine offers. Having the data in a properly built database will allow you to start using it in a creative way. For example,

with the data in a proper database, it becomes possible to create a custom data visualization that can help you understand what's going on inside your growing environment.

The images in this post show you some examples of this. The first one shows a simple example where a rather noisy humidity sensor is smoothed out using different moving averages, these averages can then be used to implement more effective control algorithms. The second image shows a detailed map of the temperature and humidity values experienced in a room, colored by the hour of the day. We can use this plot to easily locate where problematic times and VPD conditions might be, just by looking at when extreme readings happen. This behavior would be harder to observe and diagnose on a regular VPD Vs Time plot. Regular data logging web interfaces and platforms will not allow you to create plots of this sort, which is why putting your data into a proper DB and manipulating it to create custom visualizations can be very powerful.



Relative Humidity Vs Temperature map colored as a function of the hour of the day for a growing room being constantly monitored

The most powerful uses of the data come into play when you actually piece together your control and sensor data. Say you have an AC system coupled with a temperature sensor but you have a lot of other temperature and humidity readings and you also know the age of your plants at each point in time. Using this, you can create an advanced control algorithm where a system will use all of this additional information to know when to trigger AC systems and dehumidifiers to control the environment. Having a lot of logged data from a set point control system is a great starting point to train a reinforcement learning algorithm for climate control, since we know which control actions were taken at each point in time and we know the effect these had. Implementing such control mechanisms can lead to control systems that avoid spikes in

humidity and temperature across light on/off cycles, greatly smoothing out the environmental transitions for your crops.

Finally, there is also the potential to improve yields by gathering detailed mappings of yield data in a room and relating these yields with environmental sensors. If you have several different sensors in a room and you know the yield that you obtained on a per-plant basis, then you can create a map of all the yields in a room in order to see if there are important disparities in your yields because of differences in local humidity, temperature, light or air circulation levels. This can lead to important insights that can help better adjust climate conditions for the entire grow room. If multiple rooms are available, the information about environmental sensor data can be related to yields in order to stir all rooms towards more favorable conditions.

For example, after analyzing yield and temperature data from multiple growing cycles of one of my customers, we realized that the greenhouse with the lowest temperature standard deviations between sensors was giving the best yields, we then implemented better control algorithms on the other greenhouses to prevent this from happening, obtaining significantly better results across the board after that.

Data is a treasure. If you have been recording judicious sensor, climate control, and yield data through time, you're probably sitting on a gold mine that you haven't exploited yet. If you're interested in using my help to do so, please consider [booking an hour of consultation](#) time with me so that we can discuss your needs and how we could leverage your data to improve your growing results.

Commercial sensor and data logging solutions for hydroponics

On a previous post, I discussed a very interesting open-source sensor/data logging alternative for Hydroponics called [MyCodo](#), which offers a lot of features and flexibility for those growers with the time and skills necessary to implement their own sensor and data logging setup. However, many growers don't have the time to do this on their own – or the time and willingness to hire someone to do it for them – and all they want is a solution that “just works” out of the box and that fits most of their data logging needs. In this post I am going to talk about three commercial solutions – in no particular order – that I've had experience with along with some of the advantages and disadvantages that each one offers you. **Note that this post has not been sponsored by any of these brands.** *The statements below represent my opinions on the matter and the facts, to the best of my knowledge. I recommend you contact each company to ask specific questions pertinent to your needs.*



[Growtronix](#). This company offers a complete solution for monitoring and automation of hydroponic crops. Their sensors are hooked through cabled connections and they support a wide array of analogue sensors, both sold by them and by third parties. As long as a sensor can work on a 3.5-5V input and give an analogue reading, it can be installed in a growtronix setup. Their web interface is user-friendly, it allows you to

view sensor readings and create control schemes using simple if logic statements. They have also shared the source code of their web interface with some of my customers in the past, so if you would like to customize things beyond their base web application, I'm sure you could figure it out if you have the time and programming skills. Growtronix support – per the experience of the customers I have you have used it – has been stellar.

There are however some downsides to using growtronix. Since everything is cabled you will need to lay cables across your rooms if you want to hook up multiple sensors within them. The system lacks support for third party i2c sensors, meaning that you can only connect analogue sensors and will miss on some interesting third-party sensor offerings. The data is also stored in a non-relational mongoDB implementation, which means that querying data and doing complicated data analysis will not be easy with them. Their control algorithm technology is also rather simple, to the best of my knowledge they do not offer more advanced control mechanisms beyond the if logic statements they allow the users to program.



[Agrowtek](#). Similar to Growtronix, they also offer a complete monitoring and automation solution for hydroponic crops. However, they offer their own touchscreen computers to connect to their sensors, dosing pumps, and relay modules, so they do not have a dedicated web interface for their sensors that is hosted on any computer but you must purchase their own. Their “GrowControl” panels will hook with normal ethernet cables to any of the sensors they offer and you will be able to program all the behavior of the sensors and the relays from these stations. Their main advantage is easy setup, everything easily hooks up and you can then program things within the GrowControl panels to fit whatever simple control needs you might have. You can probably setup 200 sensors/relays in a day to control an entire facility using this setup. Their custom

computer also gives you more stability, meaning crashes of the system are rare (according to the customers I have who have used them). From the three companies discussed in this post, this is also the only one to offer nutrient injection systems in their offering.

However, one big limitation of this company is how closed the ecosystem is. You have absolutely no ability to hook up third-party sensors and sadly their offering lacks some important and basic sensors for a medium to large scale hydroponic setup, specifically water content and water potential sensors. You are also becoming reliant on the availability of support from them and – if the company went under – it would be very hard for you to be able to fix or find replacements for their sensors or their control panels. Their control algorithms are also fairly simple and are limited to basic if-logic, similar to the Growtronix system. Data is also not logged into any database but as basic csv files, which means substantial effort will be needed to perform advanced data analysis tasks.




[SmartBeeControllers](#). This company also offers a complete automation and monitoring solution for your hydroponic crop. Their main differentiating factor relative to the last two is that sensor stations connect wirelessly to your computer, allowing you to place sensors throughout your facility without having to set up cables through the entire place. Their sensor stations can hook up to a large number of sensors so, for example, you can use a water content station to hook up six of their capacitive water content sensors. They also require a computer server with the web software to communicate with – alike Growtronix – and their software has a focus on simplicity. In this case, control options are even more limited than in other cases, with basically only simple set-point logic available to control relays (to the best of my knowledge).

The SmartBee ecosystem is also quite limited and offers no pH/EC/ORP sensors or water potential sensors (tensiometers). You have no ability to hook up third-party sensors as well, meaning you're stuck with this offering if you use them. Because of the wireless nature of communications, sensor readings and their stability can also be compromised due to excessive electromagnetic noise, which can be particularly problematic in a short room that has a lot of HPS ballasts. It is also true that in the past (2-3 years ago) their support seemed to have problems, with several complaints about their response time online. I do not know if their technical support has improved so I would advise you to seek recent opinions about it on social media if you're considering them for purchase. The people I know who used them didn't need to contact support, so I cannot comment on this aspect from my customers' experience.

The above are three commercially available data logging systems for hydroponics. All of them should be easy to hook up and should provide you with basic data logging and control capabilities for your grow. In my opinion, the most complete one is Growtronix, given the ability to add third-party sensors – even if only analogue ones – and the quality of their sensors and web application software. However, if controlling the nutrient injection process electronically is important for your situation, then Agrowtek might be a better solution. None of them however provide advanced control mechanisms – like reinforcement learning-based climate control – and none of them provide access to all sensors that would be desirable, so a custom DIY setup might be best if these features are very important to you.

MyCodo: an open-source solution for control, data logging and visualization

There are sadly not a lot of open source solutions for data logging, visualization, and control that have all the features required to be decently expandable and allow for different use cases. Most open source solutions have been developed by individuals for their particular needs. The consequence of this is that the hardware is very specific and difficult to expand on and the software has been written to be hard-coded to the hardware, making true wide use by the DIY community hard. However, [MyCodo](#) – a project that was shared with me by a reader of the blog – seems to get rid of this paradigm, creating an open-source implementation that is truly expandable and that offers most of the features anyone would want in a truly flexible DIY setup. In this post, we will talk about this project, what it offers and how it could be expanded for hydroponic grows of all scales.



Sample control panel image taken from the MyCodo github project website

MyCodo is centered around a Raspberry Pi as its main computing hub. Once you install it, the project creates a web interface in the Raspberry Pi that you can use to manage inputs and the control actions that are derived from them. Most of the inputs that are supported by the MyCodo implementation are designed to be directly connected to the Raspberry Pi, such that the Pi acts both as the computing brains and the sensor/control hub for the implementation. When used in this way, only the Raspberry Pi is required, with whichever sensors and relays you want to add to it. This can already be powerful but has the problem that the Raspberry Pi is directly in the middle of the sensing/control environment and the entire implementation could be vulnerable to catastrophic failure due to interactions with the environment (say water getting on the Raspberry Pi).

Thankfully, the developer(s) of the MyCodo implementation had the vision to implement input/output options to use MQTT subscribe/publish mechanics. The MQTT protocol is a messaging

system where a device in a network can listen to or publish to different “topics”. So you can have an Arduino that publishes messages under the topic “HumiditySensor1” that contain the humidity value measured at each point in time and you can have it at the same time subscribe to a topic called “HumidityControl” and when it receives an “on” message in this topic it turns on a dehumidifier. You can then use MyCodo to “listen” to the humidity sensor messages, execute its own control algorithms on it, and publish the adequate control action to the “HumidityControl” topic whenever it thinks that a dehumidifier needs to be turned on. This is the way in which my custom-built Arduino/Raspberry PI control implementation generally works.



Another sample MyCodo panel

MyCodo, therefore, has a lot of flexibility that is not shared by any other open-source implementations, at least among the

ones I have found, for environmental control. Although there are no MQTT sensor stations implemented that I could find for the MyCodo, it should be fairly straightforward to build these sensing/control stations using Arduinos and the MQTT protocol and it should then be easy to add these stations to the MyCodo so that they can benefit from the system's control interfaces. In a system like this – with independent MQTT enabled sensor/control stations – you can control small or large facilities and not depend on the use of a single raspberry pi to do the entire setup. This means you could use the MyCodo to control different rooms and be able to have a centralized sensing setup for all your needs.

I have decided to give MyCodo a try for my latest hydroponic system. You should expect some videos about this in my youtube channel along with a github repository containing the code for the sensing and control stations that I am going to build in order to use Arduinos for turning relays on/off and send sensor readings. A Raspberry Pi will be used as the central control hub for the project, hosting the MyCodo webserver and code.


Pros and cons of building your own sensor and data logging system in hydroponics

If you've read my blog before, you know how important data logging is to having a successful hydroponic crop. Data allows you to monitor and tune the different variables in your grow, which allows you to give your plants the perfect environment through their entire growing cycle. However, deciding how to

do this is not simple, you need to decide if you're going to go with a company that sells some pre-made data-logging solution or you need to build everything yourself. In this post, I'm going to talk about several pros and cons of building your own data logging system for your hydroponic crop.

Pros

You have control over everything. The most important pro when building your own data logging solution is that you have total and absolute control over all aspects of it. If you want to support some type of sensors or have your data stored a certain way, there is nothing preventing you from doing this except your own skills and imagination. If you want to support an obscure messaging protocol, wireless transmission system, etc, it is all up to you. You won't be limited by the management decisions of an external company and you will be able to build a system that perfectly caters to your needs.



A simple plant monitoring custom built system. Read more [here](#). **You will be able to leverage low-cost hardware.** When building your own system you will be able to get all the parts

yourself. This means you will be able to substantially reduce costs. Of course, you're incurring the important cost of your time but the hardware itself will be low cost and once you implement the basic setup you will be able to connect new rooms and build new logging stations for a fraction of the cost of buying one commercially.

Take advantage of new hardware quickly. As new technologies for monitoring environmental variables are invented or the desire to control new variables comes into play, your ability to fully control your setup will allow you to take advantage of new hardware that comes into the scene while companies will usually be very slow to respond to such changes.

A much deeper understanding. When you build all the monitoring setup yourself, you will create a lot of understanding about how the sensors work, how each one of them is calibrated, how data is transmitted, stored, etc. If you build your own monitoring setup you will gain a much deeper understanding than somebody who just buys an off-the-shelf product.

No need for patchwork approaches. When you decide to get a commercial solution for data logging, one of the issues that comes along is that you will get the setup from a company that supports some types of sensors but you will often face challenges if you want a sensor outside this offering. This will usually mean buying a setup that includes that sensor from a completely different company, measuring some variables with one system and some others with another system.

Cons

No one to support it. The biggest drawback of building things yourself – or hiring someone to build a custom system for you – is that you will have no one to help you debug your system when things go wrong. You will also have limited ability to delegate this work, as your highly custom system will demand somebody with a high level of skill to become familiar with it

and operate it with the same level of proficiency as you do. A custom solution means all of this responsibility will fall on the shoulders of those who developed the system.



A custom built data logging system to read EC/pH/ORP. Read [more here](#).

Limited by your knowledge. Although it is true that you will get a pretty deep understanding of the things you decide to incorporate into your system, you will also be very limited in the design and implementation of your system because of your particular limitations as an individual. A big company that develops a data logging system will have dozens of people working on it, and all of their experience will go into the decisions that were made in the sensor and software implementations. This can mean better sensor choices are made, more robust communication protocols are used, etc.

Not built for sharing. Custom-built systems usually have the problem that they are built with poor documentation. Sharing

is normally not the priority and people will prefer to build “fast and dirty” in order to get things done. This means that the code is usually poorly commented and of a lower quality than what you get from a product that comes from a business. Although some people who build custom software that they intend to release as open-source implementations will often go to great lengths to provide great code quality this is rarely the case when the intention is not to make everything open source.

Big overhauls are a big problem. Since your custom building efforts will usually rely on one or two individuals, bad decisions that are made at the beginning of a project will carry a big toll during the entire life of the system. Poor decisions will be hard to overcome, as a lot of work will be needed to overhaul these “built from scratch” system. A big business with large teams will make fewer poor decision and those mistakes will be found out and fixed faster.

Messy hardware that often breaks easily. Due to the fact that people who build DIY implementations will go for rapid prototyping and functionality over robustness, sensor and data logging setups built in this manner will usually lack the roughness of commercial implementations. While a business dedicated to data logging wants to build systems with adequate sensor housing, and durability for transport, with customer satisfaction in mind, a person who builds this for him or herself might be ok with having a lot of exposed boards and cables. Overall DIY setups are therefore less robust, more likely to break, and more likely to suffer from electrical issues like poorly grounded circuitry.

Hopefully, the above pros and cons give you a useful idea of what you’re gaining and losing when you decide to build your own custom-built data logging system for hydroponics. While you will usually get much more flexible, lower cost, cohesive and personalized setups from custom building, this will usually come at the cost of higher support costs in time,

lower reliability, lower build quality, and compromises in quality depending on where your strengths as a builder/coder are. For small setups, it is usually a no-brainer to go with a custom setup – because of how much you learn from doing this and how much you can experiment – while for larger setups careful consideration of the above cons is important.

Standard hydroponic formulations from the scientific literature

When researchers started looking into growing plants without soil, they started to look for mixtures of nutrients that could grow plants successfully so that these formulations could be used to study other aspects of plant physiology. If you have a mixture of nutrients that you know grows a plant without major issues, then you can use that as a base to study other things, for example how plants react to some exogenous agent or how changes to temperature or humidity affect the uptake of certain nutrients (see this paper for a view into the history of hydroponics and standard solutions). The establishment of these standard solutions was one of the great achievements of botanists during the twentieth century, which allowed thousands of detailed studies on plants to be carried out. In this post, we're going to be talking about these standard solutions and why they are a great place to start for anybody seeking to formulate their own nutrients.


ppm (mg/L)	1	2	3	4	5	6	7	8	9	10	11	12
K	132.93	187.28	241.24	312.79	236.15	237.33	89.54	157.57	261.57	302.23	430.08	312.79
Ca	136.27	36.07	149.09	163.52	200.39	160.31	161.11	120.23	184.76	172.34	220.43	160.31

Mg	19.69	18.71	37.19	49.34	48.61	24.31	55.90	48.61	49.10	50.55	36.46	34.03
N as NH4+	0.00	4.90	2.10	18.91	0.00	28.01	19.61	0.00	0.00	0.03	0.01	17.51
Na	0.00	0.23	1.15	0.46	0.00	0.46	0.00	2.07	0.46	0.69	8.74	0.69
Fe	36.86	2.79	4.02	0.00	1.44	1.12	1.12	5.03	1.34	1.90	7.10	0.84
Mn	0.00	0.62	1.23	0.00	0.50	0.11	0.14	0.40	0.62	1.98	2.40	0.55
Cu	0.00	0.06	0.01	0.00	0.02	0.03	0.00	0.02	0.01	0.10	0.04	0.04
Zn	0.00	0.01	0.01	0.00	0.05	0.13	0.13	0.05	0.11	0.10	0.12	0.03
N as NO3	123.82	77.46	161.50	226.63	210.10	196.09	112.75	112.05	167.80	201.28	241.62	224.11
P	103.45	42.74	64.74	40.89	30.97	61.95	71.24	61.95	30.66	59.78	69.69	38.72
S	25.97	27.90	54.51	65.09	64.13	32.07	96.84	64.13	111.59	67.98	87.22	44.89
Cl	0.00	0.00	0.00	0.00	0.64	1.77	0.00	0.53	0.00	0.00	13.47	0.00
B	0.00	0.28	1.19	0.00	0.46	0.27	0.10	0.40	0.43	0.30	0.34	0.27
Mo	0.00	0.41	0.00	0.00	0.01	0.05	0.00	0.03	0.05	0.19	0.06	0.34

Summary of standard nutrient formulations found in [this article](#) with the concentrations translated to ppm. The numbers in the list correspond to the following: 1. Knop, 2. Penningsfeld North Africa, 3. Pennings-Feld Carnations, 4. Gravel Culture Japan, 5. Arnon and Hoagland 1940, 6. Dennisch R. Hoagland USA, 7 Shive and Robbins 1942, 8. Hacsalyo 1961, 9. Steiner 1961, 10. Cooper 1979, 11 Research Centre Soil-less culture, 12. Naaldwijk cucumber.

One of the best places to find a comparison between these standard solutions is [this paper](#). In it, the authors explore the relationships between the different solutions and how they are similar or diverge. In the table above, you can see a summary of the elemental nutrient concentrations found in this paper for the 12 standard solutions they compare (the paper states them in mmol/L but I have changed them to ppm as these are more commonly used units in the field nowadays). As you can see, some of the older solutions miss some elements or contain much smaller amounts of them – as they were likely present in the media or other salts as impurities – while more recent standard solutions do contain all the elements we now understand are necessary for plant life.

NH₄⁺ proportion:
 2: 0.04
 3: 0.01
 4: 0.06
 6: 0.11
 7: 0.08
 12: 0.57




Standard solutions:

1. Knop
2. Penningsfeld N. Africa
3. Penningsfeld Carnations
4. Gravel culture, Japan
5. Arnon and Hoagland
6. Dennish R. Hoagland
7. Shive and Robbins
8. Hacskalyo
9. Steiner
10. Cooper
11. Res. Centre Soil. Cult.
12. Naaldwijk, cucumber

FIGURE 1. Cation composition of the standard solutions.

Figure showing the Ca/Mg/K ratio represented in a three axis plot. Taken from the paper mentioned above.



Standard solutions:

1. Knop
2. Penningsfeld N. Africa
3. Penningsfeld Carnations
4. Gravel culture, Japan
5. Arnon and Hoagland
6. Dennish R. Hoagland
7. Shive and Robbins
8. Hacskalyo
9. Steiner
10. Cooper
11. Res. Centre Soil. Cult.
12. Naaldwijk, cucumber

FIGURE 2. Anion composition of the standard solutions.

Figure showing the N/S/P ratio represented in a three axis plot. Taken from the paper mentioned above.

It is interesting to note that all of these solutions have been successfully used to grow plants, so their convergent aspects might show us some of the basic things that plants require for growth. As they highlight on the paper, the

K/Mg/Ca ratio for most of these solutions is rather similar, as well as the N/S/P ratios. This means that most of these authors figured out that plants needed pretty specific ratios of these nutrients and these ratios are sustained with minor variations through the 12 solutions, developed across a span of more than 100 years. All the solutions developed from the 1940s have similar final concentrations and their starting pH is almost always in the 4-5 range, due to the presence of acid phosphate salts like monopotassium phosphate.

Nonetheless, there are several things that improved in the solutions as a function of time. The first is the inclusion of higher concentrations of all micronutrients with time, as macronutrient salt quality increased, the media sources became more inert and the need to add them to avoid deficiencies became apparent. The need to chelate micronutrients also became clear with time, as solutions starting with Hoagland's solution in the 1940s started using EDTA to chelate iron, to alleviate the problem of iron phosphate precipitation in hydroponic solutions. This is clearly shown in the table below, where the authors show how the first three solutions had almost or all of their Fe precipitate out, while the newest solutions, like Cooper's developed in 1979, had less than 5.5% of its Fe precipitated.

Standard solution	% Fe precipitated as Fe ₂ (PO ₄) ₃	% Cu complexed with chelate	% Zn complexed with chelate	% Mn complexed with chelate
1. Knop 1865	100	-	-	-
2. Penningsfeld North Africa	99.9	-	-	-
3. Penningsfeld Carnations	99.9	-	-	-
4. Gravel culture Japan	-	-	-	-
5. Arnon and Hoagland 1940	87.8	40.3	6.4	0.3
6. Dennisch R. Hoagland	4.0	97.5	97.5	0.1
7. Shive and Robbins 1942	99.9	-	-	-
8. Hacsalyo 1961	4.0	99.3	42.4	0.2
9. Steiner 1961	4.9	99.5	48.8	0.2
10. Cooper 1979	5.5	98.3	22.4	0.1
11. Res. Centre Soil. Cultures	6.9	100	99.2	7.7
12. Naaldwijk cucumber	4.5	96.5	7.8	0

This table shows the precipitated Fe and chelated portions of the micro nutrients in all the standard solutions.

The natural question when reading about standard solutions is: which one is the best one to use? Sadly, I don't think there's a simple answer. There have been multiple studies comparing standard solutions (see [this one](#) for an example). What ends up happening most of the time is that, while most of the solutions manage to grow healthy crops, one of the solutions happens to be more fit to the idiosyncrasies of the study because its conditions are better aligned with those that the authors developed the solutions under. A study revealing a solution to be better than another to grow plants under a given set of conditions does not imply that this solution will be the best one for all plants under all conditions. For this reason, the optimization of nutrient solutions to particular conditions using tissue analysis is still pursued in order to maximize yields.

My advice would be to view the above solutions as well researched starting points for your hydroponic crops. These solutions, especially the ones developed after 1940, will do a good basic job growing your plants. If you're interested in

making your own solutions, starting with a solution like the Hoagland, Steiner, or Cooper solutions is a great way to begin making your own nutrients. Once you have a basic standard solution working for you, you can then tweak it to maximize your yield and improve your crop's quality.

The stability of metal chelates

When you get introduced to hydroponics and nutrient solution chemistry, one of the first concepts that you learn is chelation. A chelate is a molecule formed by a metallic ion and a chelating agent – which is also referred to as a ligand – where the metal ion is wrapped around very tightly by this ligand. The job of the chelating agent is to keep the heavy metal ion shielded from the environment, allowing it to exist in solution without forming potentially insoluble compounds that will take it out of the nutrient solution. However, these chelates can be unstable or too stable, both of which can hinder the availability of the nutrient to plants. In this post, we're going to talk about what determines the stability of a metal chelate and how you can know if a given chelate will be able to fulfill its job in a hydroponic environment.



$$K_b = \frac{|ML|}{(|M| \times |L|)}$$

A simplified view of the chemical equilibrium formed $|M|$ refers to the concentration of the free metallic ion, $|L|$ the ligand concentration and $|ML|$ the chelate concentration. Charges are omitted for simplicity.

Since chelates are formed by the reaction of a metallic ion – most commonly a cation – which a ligand, a chemical equilibrium is established between the free metallic ion, the ligand, and the chelate. Every second, there are lots of chelate molecules being formed from reactions between metallic ions and ligands, and free metallic ions and ligands are being formed from the disassembly of the chelate. The process is in equilibrium when the rates of assembly and disassembly are the same. The equilibrium constant – also known as the stability constant or K_b – tells us how displaced this equilibrium is towards the product (in this case the chelate). When the K_b value is large, the concentration of the chelate at equilibrium is very large, while when K_b is small, the opposite is true. Since these numbers are usually very large for chelates, we express them as pK_b which is $-\log(K_b)$. These constants depend on temperature, but their values are independent of other chemical reactions. However, things like pH can affect the concentration of ligand or metal cation, which can affect the concentration of chelate, since the equilibrium constant's value remains the same.

	Al(III)	Ba	Ca	Co(II)	Cu	Fe(II)	Fe(III)	Hg	Mg	Mn	Ni	Sr	Zn
Acetic acid		0.39	0.53	2.24				3.7d	0.51		0.74	0.43	1.03
Adenine													
Adipic acid		1.92	2.19		3.35								
ADP		2.36	2.82	3.68	5.9				3.11	3.54	4.5	2.5	4.28
Alanine		0.8	1.24	4.82	8.18					3.24	5.96	0.73	5.16
b-Alanine					7.13						4.63		4
Albumin			2.2										
Arginine						3.2				2			
Ascorbic acid			0.19									0.35	
Asparagine			0									0.43	
Aspartic acid		1.14	1.16	5.9	8.57				2.43	3.74	7.12	1.48	2.9
ATP		3.29	3.6	4.62	6.13				4	3.98	5.02	3.03	4.25
Benzoic acid					1.6						0.9		0.9
n-Butyric acid		0.31	0.51		2.14				0.53			0.36	1

Polyphosphate			3		3.5	3			3.2	5.5	3		2.5
Proline						4.07	10			3.34			
Propionic acid		0.34	0.5		2.2		3.45		0.54			0.43	1.01
Purine					6.9						4.88		
Pyrophosphate			5		6.7		22.2		5.7		5.8		8.7
Pyruvic acid			0.8		2.2								
Riboflavin				3.9	<6					3.4	4.1		<4
Salicylaldehyde				4.67	7.4	4.22	8.7		3.69	3.73	5.22		4.5
Salicylic acid	14.11			6.72	10.6	6.55	16.35		4.7	2.7	6.95		6.85
Sarcosine				4.34	7.83	3.52	9.7				5.41		
Serine			1.43			3.43	9.2				5.44		
Succinic acid		1.57	1.2	2.08	3.3		7.49		1.2	2.11	2.36	0.9	1.78
(+)-Tartaric acid		1.95	1.8		3.2		7.49		1.36		3.78	1.94	2.68
Tetrametaphosphate		4.9	5.2		3.18				5.17		4.95	2.8	
Threonine						3.3	8.6						
Trimetaphosphate			2.5		1.55				1.11	3.57	3.22	1.95	
Triphosphate		6.3	6.5		9.8				5.8			3.8	9.7
Tryptophan							9						
Uridine diphosphate									3.17				
Uridine triphosphate			3.71	4.55					4.02	4.78			
n-Valeric acid		0.2	0.3		2.12								
Valine					7.92	3.39	9.6			2.84	5.37		5
Xanthosine				2.8	3.4	<2					3		2.4

This table was originally present in a website that no longer exists. The data is taken from the [NIST reference of heavy metal complexes](#).

The table above shows you the pKb values for different metal ions and different ligands or chelating agents. Since the pKb scale is logarithmic, a difference of 1 indicates an order of magnitude higher stability. You can also find additional references to other stability constants in this link. These constants allow us to predict which chelates will be formed if different metallic cations and ligands are present. Let's say we have a solution that contains Ca^{2+} and Fe^{3+} and we add a small amount of sodium citrate, what will happen? Since the constant for Ca^{2+} is 3.5 but that of Fe^{3+} is 11.85, citrate will chelate around 1 billion Fe^{3+} ions for every Ca^{2+} ion it chelates. In practice, this means that all the Fe^{3+} that can be chelated will be, while Ca^{2+} will remain as a free metallic ion. However, if we have Fe^{2+} instead of Fe^{3+} then Fe^{2+} has a constant of only 3.2, which means that one molecule of Fe^{2+}

will be chelated for every 3 of Ca^{2+} , meaning we will have around 25% of all the chelate formed as a chelate formed by Fe^{2+} and 75% as a chelate formed by Ca^{2+} .

We can see in this manner how chelating only one heavy metal can lead to problems. Imagine that you purchase Iron EDTA and add it to your nutrient solution, but you have added Manganese from Manganese sulfate. Upon addition, the FeEDTA chelate will disassemble to generate as much Fe^{2+} and free EDTA as dictated by the equilibrium constant and the free EDTA will then get into equilibria with all the other heavy metals, since the constant with Mn is 13.6 and that of Fe is 14.3 the ligand will redistribute itself so that it complies with all the chemical equilibria present. This means that for every 7 Fe^{2+} cations that are chelated we will have around 1 Mn^{2+} containing chelate, so you will lose around 14% of the chelated Fe in order to chelate free Manganese. That free Fe^{2+} will be unstable and precipitate out, which will shift the equilibrium and cause us to lose more of the Fe chelate. This is how competing equilibria can lead to the slow but sure depletion of available cations in solution.

With the above references and charts, you should now be able to look into any chelating agent you want to use and determine how good of a choice it is for your solution and what is likely to happen once you put that chelate in. The ligand will chelate different metals in order to comply with all the equilibrium constants, so it is up to you to add enough so that all heavy metals are satisfied or add ligands whose affinity for a given ion is so high that the others are just unable to compete for it, almost regardless of their concentration.

Six things to look for in a Hydroponic sensor data logging system

Data is key. It will help you obtain high yields and improve with each additional crop cycle. Having sensor measurements not only allows you to diagnose your crop at any given point in time but also allows you to go back and figure out what might have happened if something went wrong. With all the commercial offerings now becoming available, it is starting to become harder and harder to evaluate which data logging system might be ideal for you. In this post, I seek to share with you 5 things that I always look for when evaluating data logging systems for a greenhouse or grow room. These are all things that will enable you to store sensor data adequately and take full advantage of it, ensuring you're not handily capped by a poor starting choice.

Sensor compatibility. One of the first things that I look for is which sensors I can add and what restrictions I might have on sensors that are added to the system. I like to have systems where I can connect any 3-5V analog sensor I want. I also want to be able to connect sensors that use common protocols, like i2c sensors. I also like to know that for things like pH and EC, the boards have standard plugs I can connect to, to make sure I can replace the electrodes given to me by the company with others if I wish to do so. Freedom in sensor compatibility and in the ability to replace sensors with sensors from outside the company are both a must for me.

Expandability. Many of the commercially available data logging platforms are very restricted and can often only accommodate a

very small number of sensors. Whenever you're looking for a data logging solution that will need to be deployed on a medium/large scale, it is important to consider how this implementation can expand, and how painful it would be to make that expansion. Being able to easily add/remove sensors to a platform is key to having a flexible and robust data logging solution.



Not cloud reliant. It is very important for me to be able to use the system, regardless of whether the computers are online or not, and to have all the data that I register logged locally in some manner. Systems where an internet connection is needed for data logging or where data is not stored locally are both big show stoppers when it comes to evaluating a data logging system. There is nothing wrong with having data backed up to the cloud – this is indeed very desirable – but I want to ensure that I have a local copy of my data that I always rely on and that logging of data won't be stopped because there is some internet connection issue. Also bear in mind that if your sensors are cloud reliant you will be left without any sort of data logging system if the company goes under and those servers cease to exist.

Connectivity of sensors is robust. In many of the more trendier new systems sensor connectivity is wireless. This can be perfectly fine if it is built robustly enough, but it is

often the case that connections based on WiFi will tend to fail under environments that are filled with electromagnetic noise, such as when you have a lot of HPS ballasts. It is therefore important to consider that if you have such an environment, having most of your sensors connected using cables, or using a wireless implementation robust to this type of noise is necessary.


Have a robust API to directly access your data. Since I do a lot of data analyses using the data from hydroponics crops, I find it very crippling to be limited by some web interface that only allows me to look at data in some very limited ways. I want any data logging system I use to allow me to use an API to get direct access to the data so that I can implement a data structure and analysis the way I see fit. Having your data available through a robust API will allow you to expand the usage of your data significantly and it will also ensure you can backup your data or structure the database in whatever way you see fit. An example of this is sensor calibration logging and comparisons, while commercial platforms almost never have this functionality, having an API allows me to download the data and compare sensor readings between each other to figure out if some sensors have lost calibration or make sure to schedule their calibration if they haven't been calibrated for a long time.

Ability to repair. When making a data logging choice, we are making a bet on a particular company to continue existing and supporting their products in the long term. However, this is often not the case and we do not want to be left with a completely obsolete system if a company goes under and ceases to support the product they made. I always like to ensure that the systems that are being bought can continue working if the company goes under and that there is a realistic ability to find parts and replace sections of those products that might fail in the future if this were to be the case. *Open source products are the most ideal because of this fact.*

These are some of my top six priorities whenever I evaluate a commercial data logging solution for deployment. From the above, not being cloud reliant and having a robust API are the most important, while sensor compatibility can be ignored to an extent if the system is only being deployed for a very specific need (for which the sensors provided/available are just fine). Which of the above you give the most priority to depends on how much money you're going to be investing and how big and robust you want the implementation to be.

Differences between labels and actual composition values in commercial hydroponic fertilizers

Whenever I am hired to duplicate a company's fertilizer regime based on commercial products, I always emphasize that I cannot use the labels of the products as a reference because of how misleading these labels can be. A fertilizer company only needs to tell you the minimum amount of each element it guarantees there is in the product, but it does not have to tell you the exact amount. For example, a company might tell you their fertilizer is 2% N, while it is in reality 3%. If you tried to reproduce the formulation by what's on the label you would end up with substantially less N, which would make your mix perform very differently. This is why lab analysis of the actual bottles is necessary to determine what needs to be done to reproduce the formulations.




Average deviation from the reported composition on the label compared with lab analysis.

How bad is this problem though? Are companies just under-reporting by 1-5% in order to ensure they are always compliant with the minimum guaranteed amount accounting for manufacturing errors or are they underreporting substantially in order to ensure all reverse engineering attempts based on the labels fail miserably? I have a lot of information about this from my experience with customers – which is why I know the problem is pretty bad – but I am not able to publicly share any of it, as these lab tests are under non-disclosure agreements with them. However, I recently found a website from the Oregon government (see [here](#)), where they share all the chemical analysis of fertilizers they have done in the past as well as whatever is claimed on labels.


The Oregon database is available in pdf form, reason why I had to develop a couple of custom programming tools to process all the information and put it into a readable database. So far I

have only processed the fertilizers that were registered in 2015, but I am going to process all the fertilizers available in their database up until 2018 (the last year when this report was uploaded). However, you can already see patterns emerging for just the 2015 data. That year there were 245 fertilizers tested, from which 213 contained N, P, K, Ca, S or Mg. If we compare the lab results for these elements with the results from the lab analysis, we can calculate the average deviation for them, which you can see above. As you can see, companies will include, on average, 20%+ of what the labels say they contain. This is way more of a deviation than what you would expect to cover manufacturing variations (which are expected to be <10% in a well-designed process) so this is definitely an effort to prevent reverse engineering.



Median divergence between compositions derived from labels and

lab analyses.



Boxplot of the divergences between compositions derived from labels and lab analyses.

Furthermore, the deviations are by no means homogeneous in the database. The above graphs showing the box plot and median deviation values, show us that most people will actually be deviated by less than 5% from their label requirements, but others will be very largely deviated, with errors that can be in the 100%+ deviation from their reported concentration. In many cases, companies also have negative deviations, which implies that the variance of their manufacturing process was either unaccounted for or there was a big issue in the manufacturing process (for example they forgot to add the chemical containing the element). These people would be in violation of the guaranteed analysis rules and would be fined and their product registrations could be removed.

With this information, we can say that most people try to report things within what would be considered reasonable if the label is to remain accurate (deviations in the 1-5% range) to account for their manufacturing issues but many companies will choose to drift heavily for this and report values that are completely misleading relative to the labels. These companies are often the ones that are most widely used as they are the ones who want to protect themselves from reverse engineering most aggressively.

Take for example General Hydroponics (GH). *Their FloraGro product is registered with an available phosphate of 1%, while the actual value in the product is 1.3%, this is a 30% deviation, far above the median of the industry.* They will also not just underreport everything by the same amount – because then your formulation would perfectly match when you matched their target EC – but they will heavily underreport some elements and be accurate for others. In this same Floragro product, the K₂O is labeled as 6% and the lab analysis is 5.9%, meaning that they reported the value of K pretty accurately. However, by underreporting some but not others, they guarantee that you will skew your elemental ratios by a big margin if you try to reverse engineer the label, which will make your nutrients work very differently compared to their bottles.

As you can see, you just cannot trust fertilizer labels. Although most of the smaller companies will seek to provide accurate labels within what is possible due to manufacturing differences, big companies will often engineer their reporting to make it as hard as possible for reverse engineering of the labels to be an effective tactic to copy them. *If you want to ever copy a commercial nutrient formulation, make sure you perform a lab analysis so that you know what you will be copying and never, ever, rely solely on the labels.* I will continue working on this dataset, adding the remaining fertilizers, and I will expand my analyses to include

micronutrients, which are covered by Oregon government tests.