76th International Astronautical Congress (IAC 2025), Sydney, Australia, 29 Sep-3 Oct 2025.


IAC-25-A6.7.1 (99453)

# Understanding and Mitigating Data Limitations for Conjunction Assessments

**T.S. Kelso[a]*, Kevin Kuciapinski[a]**

[a] *CelesTrak, 640 Luakalae Place, Wailuku, HI, USA 96793*
* Corresponding Author

## Abstract

Increasing orbital congestion, with satellites that maneuver more frequently, has uncovered serious weaknesses in current conjunction assessment (CA) processes that pose increased risk to safety of flight, if not fully understood.

Historically, space operators have accepted, used, and trusted orbital data and conjunction assessment products produced by the US government, without question, and many users are left with the false impression that they have it all under control.

Yet these products can be readily shown to have serious flaws by using more reliable orbital data products like satellite operator ephemerides or other high-precision orbital data products. It is imperative that operators understand how to easily use existing tools to detect when these flaws are present and apply available data to quickly remedy them.

There are also flaws in standard data products, like the Conjunction Data Message (CDM), that undermine the ability to screen with these more reliable orbital data products because they actually mask the effects of maneuvers due to improper force modeling. Obviously, this has serious implications for constellations, like Starlink, where satellites are maneuvering all the time.

The good news is that these existing solutions are already available on CelesTrak and can quickly help operators not only identify bad data but actually provide data to resolve problems like missed maneuvers and cross-tagging.

This paper will examine the conditions and faulty assumptions that lead to this situation and show how the tools to address them were developed and how they can be used by SSA providers to improve overall SSA. It will also show how a simple—and obvious—change to the CDM can remove a fundamental flaw. Finally, this paper will address common misperceptions about how other data products, such as SP (Special Perturbations) or the new XP data type can mislead operators into a false sense of security, when they also suffer the same fundamental issues as GP (General Perturbations) data when working with maneuvering satellites.

Ultimately, it is hoped that a more complete understanding of these issues and remedies will encourage satellite operators and SSA providers to work together more closely to apply these lessons-learned and motivate others to freely and transparently share their data to ensure we are best prepared to do what is necessary to protect safety of flight and protect the near-Earth orbital environment.

**Keywords: Conjunction Assessment, Orbital Data, Orbit Propagators**

## 1. Introduction

Whether satellite operators want to admit it or not, a lot of other operators are still using GP data for CA. There are a number of reasons for this. These include data size, readily available software using validated versions of SGP4, and computational complexity. In addition, while ephemerides may be available for some active satellites, they aren't for dead satellites, rocket bodies, or other debris. And not all satellite operators make their ephemerides readily available.

In those cases where ephemerides are available, the sheer amount of data needed can be overwhelming—even before you consider how you are going to store and process all of it. In fact, CelesTrak currently spends about 1 hour 45 minutes per Starlink ephemeris updates (typically 3 times per day) for ~8,400 ephemerides totaling ~17 GB of data. That's an hour and 45 minutes of extra data latency from the data creation time. If someone is uploading that to some data repository and then an operator is downloading it, that's twice the data latency increase. That doesn't include any CA processing time.

So, what's the real problem with using GP data, anyway? Sure, the force models are pretty simplistic and they don't include maneuvering. They don't work well for high-eccentricity orbits, since SGP4 was developed using series expansions in eccentricity. They are limited to Earth orbits (not a problem for most of us yet). And they don't include covariance. But what about accuracy?

We're all told they aren't very accurate, right? Well, we are told that, but do you understand why? It might not be what you think. If you think it's because of the SGP4 force model, you would be wrong. Studies have

shown [1] that it can be quite accurate using good data. But there are several other issues causing inaccuracies.

First, the Space Surveillance Network (SSN) sensors still only send back 3 observations per pass—just like they have for many decades now. A first, middle, and end point. It doesn't make a difference which sensor or its capabilities. Then, most SSN sensors are in the Northern Hemisphere, so orbits may be good over that part of the orbit, but we all know conjunctions occur equally over both hemispheres.

And then there is fit span. If you choose a small fit span, you can get smaller residuals, but the quality of the orbit in the Southern Hemisphere will be worse. So, why would you want to use a smaller fit span? To avoid processing through unknown—and unmodeled—maneuvers.

Current maneuver-detection techniques within the 18th Space Defense Squadron (18 SDS) aren't particularly efficient at detecting maneuvers. Back when these techniques were first developed, there weren't very many satellites and satellites didn't maneuver very often. Clearly that has changed significantly.

And before you start thinking, "But there is Special Perturbations (SP) data that solves all of this," I'm going to ask you if you *know* that or simply heard that. Yes, SP data is generated from numerically integrated orbits that can have much more sophisticated force models than SGP4. But it still uses the same SSN data, with the same SSN sensor placement limitations. And, do you know what isn't in their force modeling: maneuvers. Missing maneuvers in the SP data is just as bad as doing it in the GP data.

Now you might still believe the SP data is somehow much better than GP data, but let’s correct that assumption now. It turns out that 18 SDS produces the GP data *from* the SP data using a process called eGP, which fits an SP ephemeris using SGP4 to produce GP data. So, unless 18 SDS is somehow misapplying eGP, the GP data will not be fundamentally much worse than the SP data (more on that later). And, turning that around, if the GP data is bad, so is the SP data (more on that later, too).

These problems with the SSN and 18 SDS processes aren't due to the people using them. They are due to a failure of US Air Force (USAF) and now US Space Force (USSF) leadership to anticipate future needs and adequately resource, develop, and implement solutions. Because of these failures, at least four upgrades of these systems have failed and USSF is stuck with hardware and software that hasn't changed since the 1980s.

On top of that, they rely on non-cooperative tracking of the entire satellite catalog—because back at the beginning, the US knew they couldn't depend on the Soviet Union to help. But that applies to all US satellites, as well, including those operated by the US Department of Defense (DOD). And the one thing we know about non-cooperative tracking is that the hardest objects to track are the ones that maneuver. And before you ask, no 18 SDS does not use the ephemerides uploaded to them via Space Track for Space Situational Awareness (SSA) (we will show this soon).

Sometimes you really do need to pay attention to history to understand how we got where we are.

## 2. The Development of SOCRATES (2004)

So, it seems we are doomed from a data quality perspective, right? The most comprehensive source of orbital data has a variety of issues that are due to limitations resulting from being stuck with legacy software from the 1980s, inadequate sensor placement, lack of force modeling for maneuvers, and a dependency on non-cooperative tracking.

But what if we rethink this problem? The good news is that CelesTrak started down this road just over twenty years ago when we first developed SOCRATES to show it was possible to screen all payloads against everything, using COTS hardware and software and standard orbital data (GP data), do that within a reasonable amount of time, and make the results readily available to all users via the Internet.

This demonstration was intended to motivate the community to show that serious conjunctions were not a rare occurrence and to raise awareness that we needed to do more—and that we could. It also began to reveal the real issue with the process—that the GP data wasn't suited to the task of performing CA. After all, SGP4 and the GP data were designed to be good enough to maintain track custody as objects passed from sensor to sensor, to allow more efficient use of those sensors (tracking an expected object as opposed to searching for it). They weren't designed for CA, back in the days when everyone believed 'space is big.'

We did a pretty good job of convincing the community of those facts. But that didn't improve anything. There was lots of talk about more sensors and better processing hardware (requiring substantial planning and investment). But that would take a long time. What could we do now?

We already knew that the hardest things to track were active satellites that maneuvered to maintain their mission requirements. So, we decided, why not just ask satellite operators for their ephemerides, which they already produced for their own use, and use that instead of GP data when available.

We started asking in Jul 2007 with Iridium (unsuccessful) and just after that, Intelsat approached us and asked, "What if we used our ephemerides instead of GP data inside of SOCRATES?" That soon grew to include SES, Inmarsat, Telesat, and Eutelsat and by the end of the year, SOCRATES-GEO was born. This would ultimately lead to the formation of the Space Data Association (SDA) and the Space Data Center (SDC). It also led to increasing usage of ephemerides in CA screening. And the SDC made it possible to directly share conjunction results between parties and they could immediately reach out by e-mail or voice using shared contact information.

Just after trying with Iridium, in Aug 2007, we also approached JSpOC about adapting SOCRATES to use their SP data to allow them to screen more satellites and do that much more efficiently (unsuccessful) inside JSpOC.

But we still had this fundamental misalignment of results when USAF wasn't yet using operator-provided ephemerides and relied on SP data instead, primarily due to missed maneuvers. Even now, when 18 SDS uses operator ephemerides, it's not clear which results use operator ephemerides and which use SP. And we still had the issue that some operators were relying on GP data alone. It was clear that we could still do better.

## 3. Realizing the Need for Supplemental GP Data (2007)

At the end of 2007, I received an e-mail from a CelesTrak user in Brazil, reporting that the position calculated by SGP4 using the latest GP data for one of the active GPS satellites was off by over 20,000 km from the position calculated for that object using data from the GPS almanac. How could that be? He asked if I knew which one was right. I had actually done a detailed paper assessing accuracy of the GP data for GPS satellites in January of that year—"Validation of SGP4 and IS-GPS-200D Against GPS Precision Ephemerides"—so I had a pretty good idea what the answer was.

You can read more about that story on CelesTrak in my SDA TAP Labs talk from Feb 2025, but it quickly became clear that the GP data hadn't been updated because the GPS satellite had maneuvered and JSpOC lost track custody. Despite JSpOC and the 2nd Satellite Operations Squadron (2SOPS) both being part of the USAF and 2SOPS putting out advance notification of maneuvers to the world (via NANUs) or that they shared high-accuracy orbital data in the form of almanacs that were used by every GPS receiver on the planet, there were no communications between these organizations and JSpOC could only use GP data (in the TLE format) at the SSN sensors or by their analysts. So, why not fit the GPS almanacs, propagated IAW the GPS ICD, and then fit that data with SGP4 to produce Supplemental GP (SupGP) data?

That is, instead of using non-cooperative tracking with no positive ID and subject to issues with everything from fit spans to sensor placements—and both recent and planned maneuvers—use the almanac data instead. When comparing the GP and SupGP data for GPS—and then GLONASS—it was clear the new SupGP data was clearly superior. We had found a way to continue to use a very computationally efficient propagator (SGP4) and data efficient storage (TLE-formatted GP data) to allow anyone using SGP4-compatible software to incorporate high-accuracy results—including JSpOC and now 18 SDS. You can see the results of those early validation tests on the Supplemental GP Element Sets page on CelesTrak—right under links to all the sources of independent operator data we use today to generate SupGP data for 9,621 satellites of the 12,557 active satellites that have 18 SDS GP data (76.6%) [as of 2025 Sep 12].

## 4. Putting It All Together

Those first SupGP data sets were produced almost 18 years ago, but the problems haven't gone away. In fact, we still see GPS satellites go untracked for days following announcements in the GPS NANUs of 'forecast delta-Vs." As of 2025 Sep 12, the GP data for 62339/GPS BIII-7 is now ~11 days old and there is an RMS difference between the SupGP and GP data of 89.269 km. This despite me sending many detailed e-mails to 18 SDS showing them how easy it is to use CelesTrak to not only see when these problems exist, but to just as easily fix them.

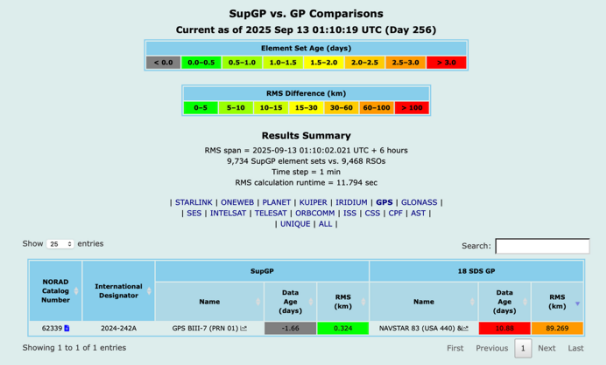


**Figure 1. SupGP vs. GP for GPS B3-7**

So, if 18 SDS isn't putting the SupGP data to use to find and fix errors in their GP data (and, as we saw earlier, in the source SP data), the onus is on satellite operators to make sure they do it themselves. One of the first hints of trouble can be the age of the GP data. Since we can only fix problems for active satellites we have independent orbital data for (normally in the form of ephemerides), we set up a dynamic query that shows the oldest GP data for active satellites at:

https://celestrak.org/NORAD/elements/table.php?GROUP=active&SHOW-OPS&OLDEST

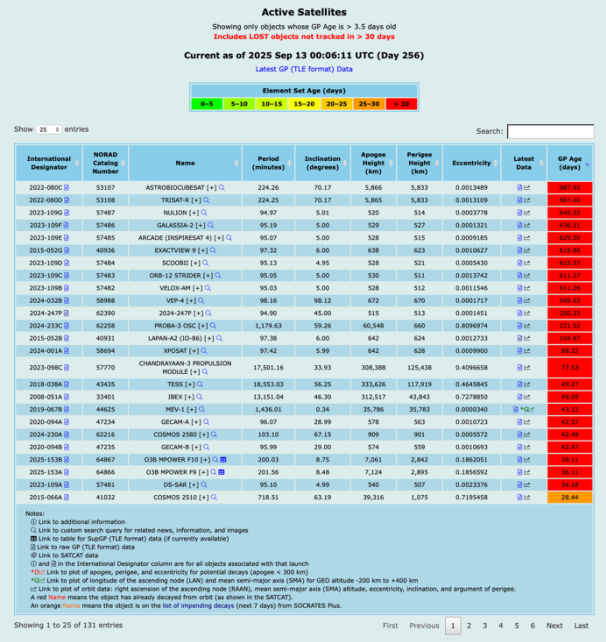


**Figure 2. Active Satellites with Oldest GP Data**

You might be surprised to learn that there are more than a few (supposedly) active satellites that have been 'lost' for more than a year. Some are in orbits that are known to be difficult to track (and that SGP4 really shouldn't be used for). But others involve satellites where their operators provide ephemerides to 18 SDS (via Space Track) for CA screening. Two of those right now are O3B MPOWER F9 and F10—both with GP data over 36 days old. Any active satellite with SupGP data will have a table icon following its name with a link to the SupGP table and the data (linked in the Latest Data column). There is also a link to a graph of the past 2.5 years of SupGP data to see whether there were any recent maneuvers.

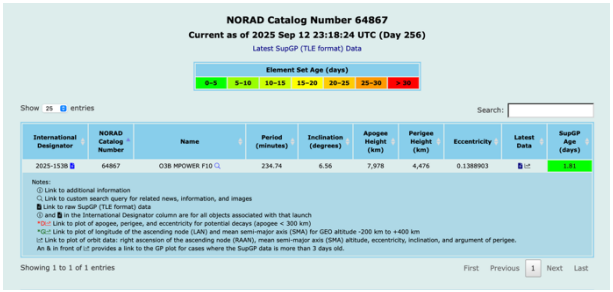


**Figure 3. SupGP Table for O3B MPOWER F10**

But GP data age isn't the best way to find problems quickly and doesn't show our work. For that, we created the SupGP vs. GP Comparisons table:

https://celestrak.org/NORAD/elements/supplemental/table-rms.php?UNIQUE&WORSE-THAN=25

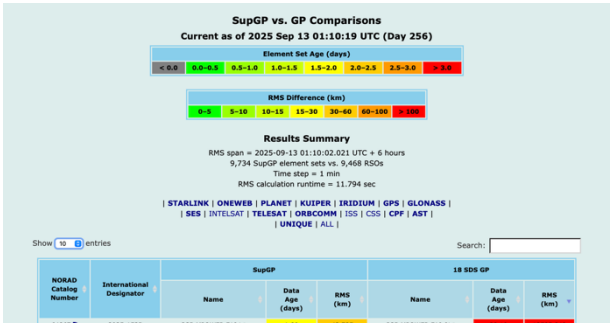


**Figure 4. SupGP vs. GP Comparisons**

There is a lot to unpack here. To begin with, the Results Summary shows the RMS span and then that we are comparing [as of 2025 Sep 12] 9,734 SupGP elements to 9,468 GP element sets. We sometimes get independent data for a satellite from multiple sources. Note particularly that this is an especially easy thing to do: To take 9,734 SupGP element sets and compare them to the corresponding data at 1-minute intervals for 6 hours only took 11.794 seconds. There is no reason not to do these calculations and we refresh them every

30 minutes. And below those results are the individual SupGP data sets.

In the table, the first column contains the NORAD Catalog Number and a link to the latest SupGP for that object, followed by the International Designator. Then there is the SupGP data for that object (not always the SATCAT name) followed by a link to a graph showing the last 2.5 years of SupGP data. Then there is the SupGP data age. The RMS is the value determined when we fit the source orbital data with SGP4 to produce the SupGP data. It's usually quite good, but can be thrown off by large maneuvers or reentry (both hard to model with SGP4). Essentially, that's how well the data fit over our 8-hour fit span. There is a similar set of data for the 18 SDS GP data, but here the RMS difference is between the SupGP and GP data. So, if the SupGP vs. source RMS is ~600 m but the GP vs. SupGP RMS is ~4,400 km (an actual case from today), it should be clear where the problem is (Answer: the GP data and likely any SP data). And each of the individual columns is sortable, so if you're curious what the worst SupGP RMS is, that's easy to see (Answer: it's the STARLINK-31204 data that's only 0.23 days old and predicted using the SupGP data to have reentered a few hours ago on 2025 Sep 12).

You can change the 'worse-than' threshold to anything from zero up. And you can look at a single result using a query like:

https://celestrak.org/NORAD/elements/supplemental/table-rms.php?CATNR=64867&WORSE-THAN=25

Looking at the Starlink data at the time Fig. 5 was captured, we see 1,031 entries with a GP RMS over 25 km, out of 8,381 with SupGP data (12.3%):

https://celestrak.org/NORAD/elements/supplemental/table-rms.php?STARLINK&WORSE-THAN=25

The oldest SupGP data is 0.63 days old or less than 8 hours. The GP data doesn't have to be much older to throw things off significantly if the satellite is maneuvering.

A quick way to see the impact is to look at our GP vs. SupGP matching page for Starlink. Here we go launch by launch and compare every SupGP element set to every GP element set looking for any match(es) within 30 km. Numbers in red show how many objects have NO MATCH. Asterisks show matches with other objects and possible cross-tags.

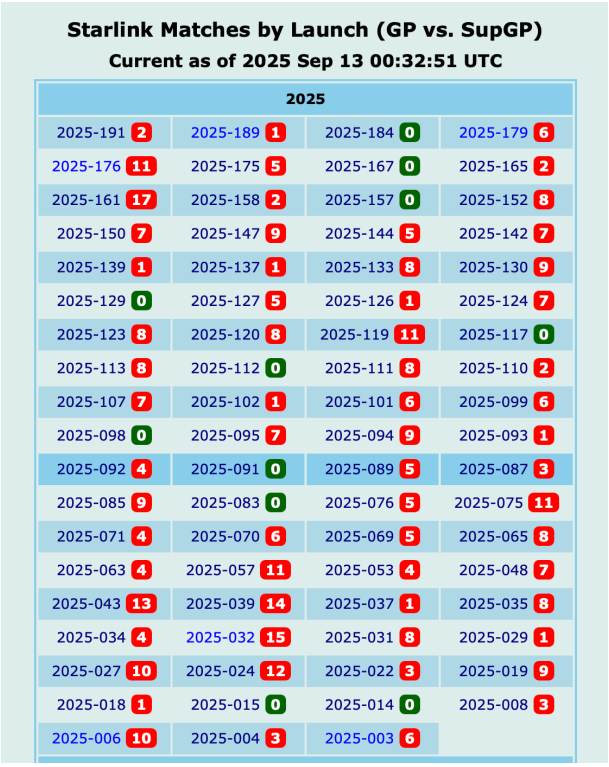

**Starlink Matches by Launch (GP vs. SupGP)**

**Current as of 2025 Sep 13 00:32:51 UTC**

| 2025 | | | |
|---|---|---|---|
| 2025-191 2 | 2025-189 1 | 2025-184 0 | 2025-179 6 |
| 2025-176 11 | 2025-175 5 | 2025-167 0 | 2025-165 2 |
| 2025-161 17 | 2025-158 2 | 2025-157 0 | 2025-152 8 |
| 2025-150 7 | 2025-147 9 | 2025-144 5 | 2025-142 7 |
| 2025-139 1 | 2025-137 1 | 2025-133 8 | 2025-130 9 |
| 2025-129 0 | 2025-127 5 | 2025-126 1 | 2025-124 7 |
| 2025-123 8 | 2025-120 8 | 2025-119 11 | 2025-117 0 |
| 2025-113 8 | 2025-112 0 | 2025-111 8 | 2025-110 2 |
| 2025-107 7 | 2025-102 1 | 2025-101 6 | 2025-099 6 |
| 2025-098 0 | 2025-095 7 | 2025-094 9 | 2025-093 1 |
| 2025-092 4 | 2025-091 0 | 2025-089 5 | 2025-087 3 |
| 2025-085 9 | 2025-083 0 | 2025-076 5 | 2025-075 11 |
| 2025-071 4 | 2025-070 6 | 2025-069 5 | 2025-065 8 |
| 2025-063 4 | 2025-057 11 | 2025-053 4 | 2025-048 7 |
| 2025-043 13 | 2025-039 14 | 2025-037 1 | 2025-035 8 |
| 2025-034 4 | 2025-032 15 | 2025-031 8 | 2025-029 1 |
| 2025-027 10 | 2025-024 12 | 2025-022 3 | 2025-019 9 |
| 2025-018 1 | 2025-015 0 | 2025-014 0 | 2025-008 3 |
| 2025-006 10 | 2025-004 3 | 2025-003 6 | |

**Figure 5. Starlink GP vs. SupGP Matching**

A lot of this is better discussed in my SDA TAP Lab talk, where we don't have to try to cram tables with lots of information into a 2-column format. I highly encourage you check that out for additional understanding of this important topic.

## 5. Implications for Conjunction Assessment and Safety of Flight

So, we know that many satellite operators upload ephemerides to Space Track every day for 18 SDS to screen for close approaches. This is the result of the work with SOCRATES back in 2008 to show that it could and should be done that way. And if you have the data used, it should be easy to verify that the calculations for a particular close approach were done properly. Of course, just because it is easy doesn't mean it shouldn't be checked—each and every time.

Many years ago, when I was in charge of SDC operations, we got *lots* of CDMs (Conjunction Data Messages) from 18 SDS for our operators. We had a pretty tight process where operators uploaded their data to us and we uploaded it to Space Track in a way we could identify whether a particular CDM was using the same ephemeris as we were. Even though we were doing everything we could to reduce the latency involved in producing results, it wasn't unusual that the CDMs we got would be using an earlier ephemeris. But

on occasion, we would find we would receive CDMs using the same ephemerides and the results were matching ours.

That's great. It worked once, so we don't need to do that any more, right? Wrong.

One day, I noticed a CDM that reported a close approach that used the same ephemeris as we did but had a dramatically different set of calculated circumstances. How could that be? I spent a lot of time trying to see what we might be doing wrong. Then, as I was looking over the raw CDM data for the second or third time, I saw it. In the optional data that reported the apogee and perigee for this GEO satellite, it showed a negative perigee. Wait, what? Yes, it showed a perigee below the surface of the Earth.

I immediately reached out to 18 SDS to try to determine how that could be—and how a negative perigee wouldn't be flagged as an error. I asked how they could interpolate the ephemeris data to come up with that result. The response I got puzzled me at first: "Oh, we don't interpolate, we fit." What does that mean? 18 SDS couldn't seem to answer that question.

Eventually it dawned on me. When the CDM format was created, it was done via the CCSDS standards process. DOD wasn't interested in going through an extended standards development and already had a process and format (well, specific data elements) to do that. DOD insisted that the community use their format. In that format, it included yet another way to represent orbital data.

At the time, I was pretty adamant that this was a mistake. CCSDS had already developed three orbital data formats: OEM (Orbit Ephemeris Message) for ephemerides, OPM (Orbit Parameter Message) for state vectors, and OMM (Orbit Mean Element Message) for GP data. I insisted that there was no reason why both objects involved in a conjunction would use the same type of orbital data and that rather than converting a standard format to the state vector being proposed for the CDM (via some unknown process), the message should include the original OEM/OPM/OMM data actually used for the calculation.

Of course, the CDM was designed to handle SP data, which was in the form of a VCM or Vector Covariance Message. At the time, JSpOC still hadn't come to accept why they should screen using operator ephemerides, as SOCRATES and SDC had already been doing.

If you look at the CDM format, you find there are two state vectors and an associated set of force models for each. And guess which force isn't included? That's right, maneuvers. There is no way to include maneuvers, as can be done with an ephemeris. Of course, to get everything to match, you have to re-fit the ephemeris with the force model specified in the CDM. That works fine if there are no maneuvers. But if there are maneuvers, you have just smoothed out any maneuver information in the ephemeris and if a close approach occurs around a maneuver, the fit will be worse.

It's still not clear what part of an operator ephemeris is being fit to produce a CDM, but the fit occurs before the calculations, to ensure they match.

This is yet one more example of people not thinking through the implications of a decision beforehand and not checking—or validating—their assumptions on each and every event. And once again, this is a situation where there is an easy solution—that was already being used in the SDC. The CDM should not specify yet another orbital data format, instead it should simply include the original data in the appropriate format—OEM, OPM, or OMM—as part of the package.

## 6. Conclusions

By now, I hope I have convinced you that blind reliance on the 18 SDS data products—whether that's GP, SP, or the (possibly) soon-to-be released XP data—is not advisable. They all suffer the same fundamental limitations and are inextricably linked. In these days of proliferated constellations that are constantly maneuvering, keeping up with non-cooperative tracking that can take days to detect and process maneuvers is not giving us a timely enough picture of what's going on.

While ephemerides can help, the sheer volume of data only increases data latency. If you're waiting for data to upload to Space Track and then plan to download it all, that further increases data latency. But converting that data to SupGP data allows users to grab a complete data set for far less than a single ephemeris file, minimizing further data latency. And the regularly updated tables on CelesTrak can help you quickly answer the question of which data you can trust if you think you have a serious conjunction. We expect to add direct links to the SupGP vs. GP Comparisons for the 2 objects involved in a conjunction in SOCRATES Plus and add graphs, as well. And we're looking into when we might use SupGP instead of GP or make it easy to switch to those results. If nothing else, it should give you better situational awareness for not just single events but the larger picture. And being informed of the real problems is when we begin to develop real solutions.

But we're not done. While having SupGP data for ~76% of the objects with GP data is a start, we can do better. CelesTrak already has access to data from Starlink, OneWeb, Kuiper, and future proliferated constellations like Telesat and AST Space Mobile, but we need *every* satellite operator to pitch in. China is launching Hulianwang Digui (Guowang) satellites regularly and with Qianfan, they are each targeting proliferated constellations of 10,000 or more satellites. We already see problems with old data for these satellites. We regularly have problems on Transporter and Bandwagon launches where objects aren't identified for months (some never). Even if their operators know where they are at, if they don't share their best data, that just increases the chances that some satellite—that they don't know its orbit well enough—could run right into them. Sharing is key here and CelesTrak is ready to work with *all* satellite operators to push that number to as close to 100% as possible.

And that raises one last issue. Trying to assess operational or ops status is an ongoing challenge. CelesTrak regularly reviews open-source materials, works with satellite operators, and even reviews orbital data to try to infer changes in ops status. So, why does that matter? Well, maneuvering satellites are dangerous enough if you have timely information on their orbit. But a dead satellite can be even more so. Take the case of Starlink, who asserts it will autonomously maneuver out of the way of any dangerous situation (we'll leave the question of what data they are using to make that determination to another time). So, if you think your satellite might be on a trajectory for a very close approach, you might just sit tight. But what if you didn't know that satellite wasn't fully operational and might have maneuvering or attitude control issues? Starlink doesn't really share that information and CelesTrak spends lots of time once a week reviewing orbital data in an attempt to keep that current. All satellite operators need to be much more transparent on the ops status of all of their satellites so that others don't misinterpret whether their satellite can or cannot move during a close approach. Transparency is key here and again CelesTrak is here ready to share that data with the rest of the space community.

*List of references*

[1] Oltrogge, D.L. and Ramrath, J., "Parametric Characterization of SGP4 Theory and TLE Positional Accuracy," AMOS 2014, Maui HI, 11 September 2014, accessible at https://comspoc.com/Resources/Content/Private/C-20220424T080117/Paper/20140911_EGP_Fit_Capability_AMOS_3317135_v04.pdf.