AI CASE STUDY: CLAUDE AND CCSD(T) DERIVATIVES

PROF. T. DANIEL CRAWFORD, VIRGINIA TECH

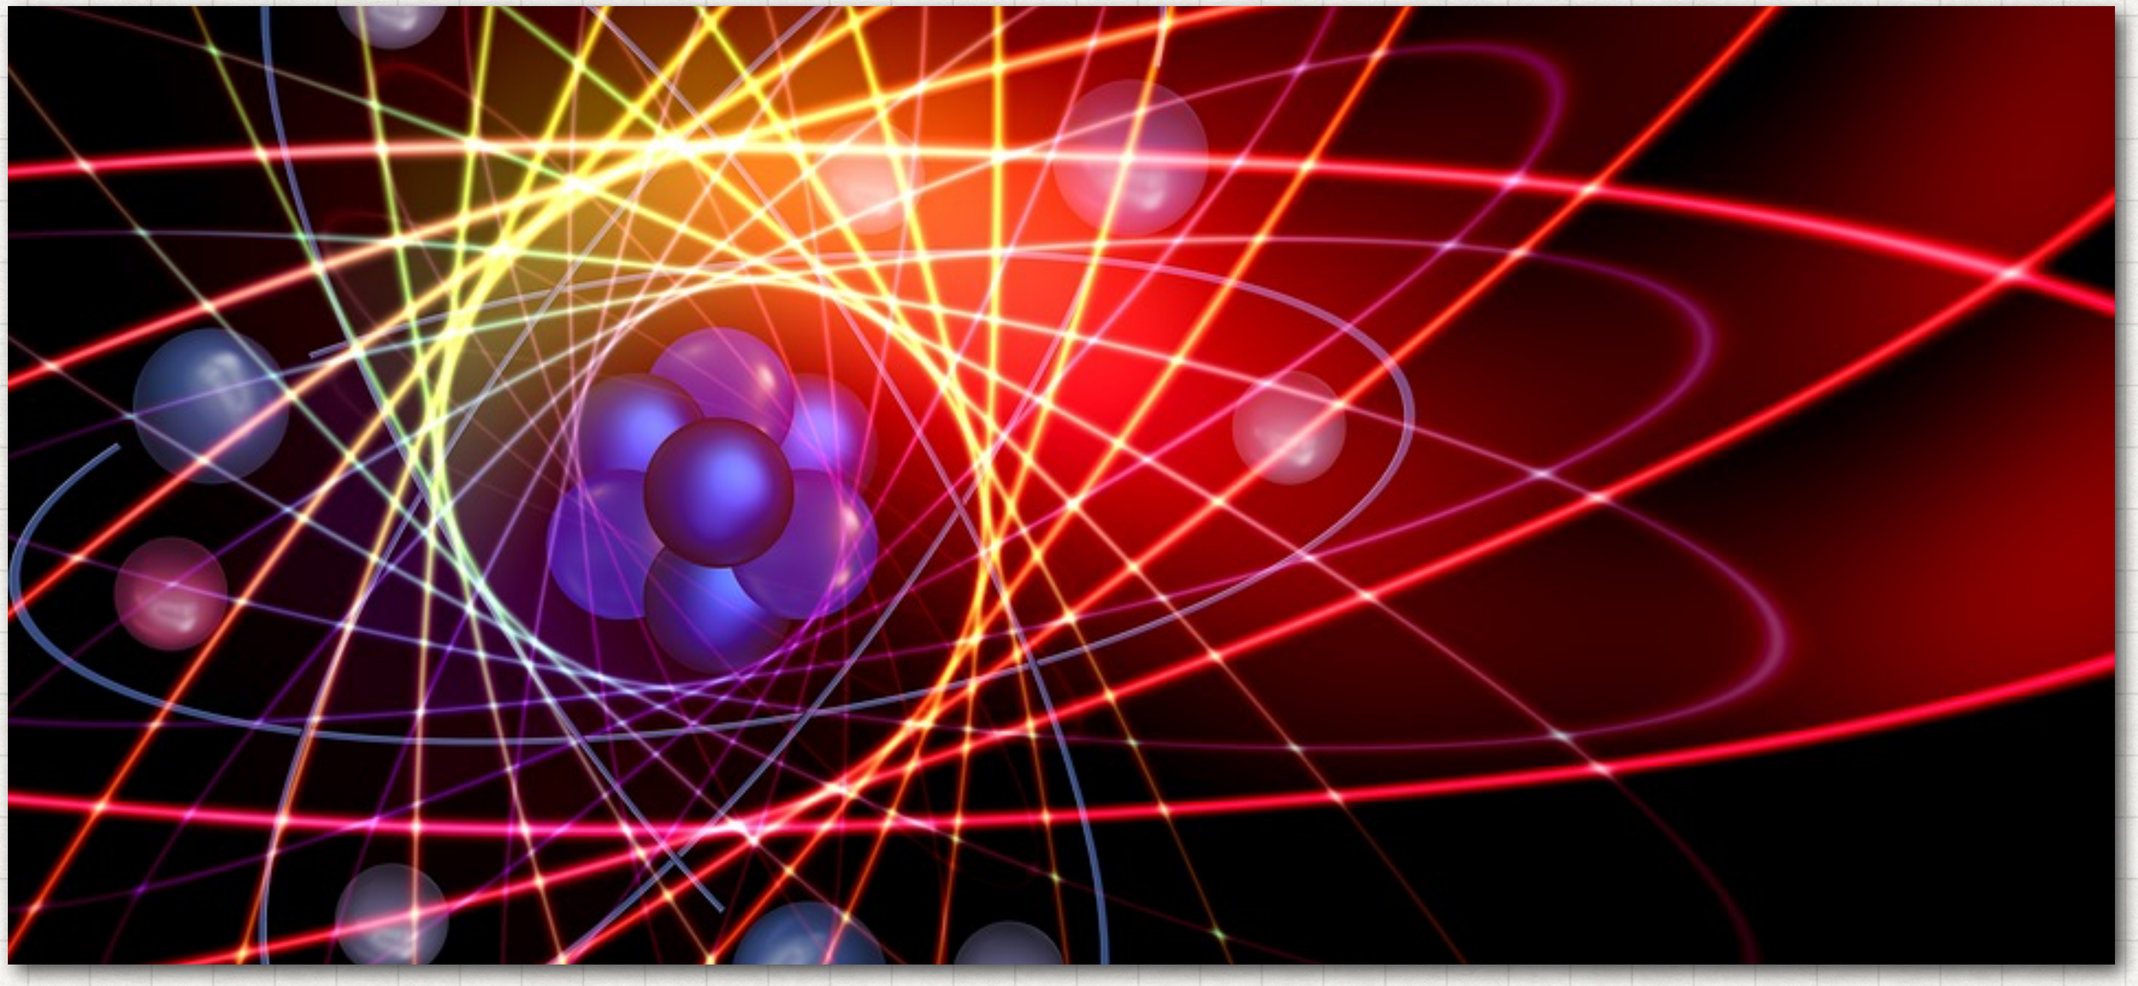

THE CONTEXT: CCSD(T) ANALYTIC DERIVATIVES IN PYCC

README Contributing BSD-3-Clause license

PyCC

License BSD 3-Clause | CI passing | codecov 87% | docs passing

A Python-based electronic-structure package centered on coupled cluster, built on a shared `Wavefunction` base that also hosts Hartree-Fock, MP2, and CI. Current capabilities include:

- RHF-/UHF-/ROHF-based CCD, CC2, CCSD, CCSD(T), and CC3 energies, via either a spin-adapted (closed-shell RHF) or a spin-orbital formalism; the latter enables the open-shell (UHF/ROHF) correlated methods
- Triples drivers for various approximate triples methods
- One- and two-electron (reduced) density matrices for CC2 and CCSD (with CC3 and (T) contributions), and CC one-electron properties (e.g. dipole moments)
- EOM-CCSD excitation energies for both the spin-adapted (closed-shell RHF) and the spin-orbital (open-shell UHF/ROHF) references
- Linear response functions (dynamic polarizability, optical rotation) for both the spin-adapted (closed-shell) and spin-orbital references
- Real-time (RT) CC2, CCSD, and CC3 with a selection of integrators
- Local (PAO, PNO, PNO++) CCSD energies, and local RT-CC
- MP2 (`MPwfn`) and CISD/CID (`CIwfn`) energies (closed- and open-shell)
- Analytic MO-basis derivative properties for Hartree-Fock (`HFwfn`) and MP2 (`MPwfn`): energy gradient, static dipole polarizability, nuclear Hessian (force-constant matrix), atomic polar tensors (APTs / dipole derivatives, in both the length and velocity gauge), and atomic axial tensors (AATs) -- the building blocks for IR and VCD spectra. Both spin paths (spin-adapted closed-shell RHF and spin-orbital), all-electron and frozen-core; the MP2 second derivatives come via two independent routes (explicit-derivative and 2n+1)
- Analytic derivative properties for CCSD and CCSD(T): nuclear energy gradients, relaxed electronic dipoles, static dipole polarizabilities, nuclear Hessians (force-constant matrices), and atomic polar tensors (APTs / dipole derivatives). Both spin paths (spin-adapted closed-shell RHF and spin-orbital UHF), all-electron and frozen-core, built on the same relaxed-density / Z-vector and 2n+1 machinery as the MP2 derivatives

https://github.com/CrawfordGroup/pycc

CASE 1: CCSD(T) POLARIZABILITIES

AN ELEGANT SYMMETRY EXPLANATION THAT WAS SIMPLY WRONG

The static polarizability is the second derivative of the energy with respect to an external static field:

$$\alpha_{ab} = -\partial^2 E / \partial F_a \, \partial F_b = \sum_{pq} \frac{\partial \tilde{D}_{pq}}{\partial F_b} \, \mu^a_{pq} + \sum_{pq} \tilde{D}_{pq} \sum_t \left(U^b_{tp} \, \mu^a_{tq} + U^b_{tq} \, \mu^a_{pt} \right)$$

which requires derivatives of the relaxed one-electron density, the first-order coupled-perturbed Hartree-Fock coefficients, and derivatives of the electric dipole integrals.

While debugging the code against a CFOUR “oracle” for CCSD(T)/6-31G for HOF, we observed errors on the order of 10^{-7}:

component	pycc	CFOUR	|pycc – CFOUR|
α_xx	-2.21066499	-2.2106648649	1.25e-07
α_xy	-0.72206744	-0.7220673958	4.42e-08
α_yy	0.16615317	0.1661531583	1.17e-08
α_zz	0.06526038	0.0652603834	3.40e-09

CASE 1: CCSD(T) POLARIZABILITIES

AN ELEGANT SYMMETRY EXPLANATION THAT WAS SIMPLY WRONG

After some discussion of the potential source of the discrepancy, the AI confidently produced an explanation:

> **The anomaly is localized to the out-of-plane (x) component.** [...]
>
> So the mechanism is real and I finally have it right, with the symmetry argument behind it:
>
> - A **planar molecule has σ_h symmetry**; MOs are a′ (in-plane-symmetric) or a″ (antisymmetric). The dependent-pair numerator l′_ij – l′_ji between an a′ and an a″ orbital is **exactly zero** at the planar geometry (different σ_h symmetry). An **out-of-plane field is an a″ operator** — it turns those numerators on from zero, so they cross the 1e-8 gate under the field. An **in-plane field preserves σ_h**, so no forbidden pair is activated → clean alpha. That's exactly the y/z-clean, x-anomalous pattern.
> - This is **precisely the case Lee-Rendell's gate is designed for** [...]
>
> **And it is not a bug:** [...] clean everywhere **except the out-of-plane direction of a planar molecule**, a σ_h + gate — not a wrong value anywhere.

This is not a hallucination, not a fabricated citation, and not obviously silly. It cites the correct paper, invokes a real selection rule, explains an earlier failed test of its own (a "tilted C_1 water" that was still planar), reconciles itself with the literature, and fits the observed x/y/z pattern quantitatively. **It is also false.**

CASE 1: CCSD(T) POLARIZABILITIES

AN ELEGANT SYMMETRY EXPLANATION THAT WAS SIMPLY WRONG

I did not agree with Claude's explanation for the discrepancy, which involved aspects specific to CCSD(T) (i.e., the use of canonical perturbed MOs – see Case 2 later) and I asked it to produce the same comparison for CCSD, which does not use the same formulation for the MO response:

component	pycc	CFOUR	|pycc – CFOUR|
α_xx	-2.2351823465	-2.2351822112	1.35e-07
α_xy	-0.7098118054	-0.7098117576	4.78e-08
α_yy	0.1333827886	0.1333827763	1.23e-08
α_zz	0.0617969674	0.0617969710	3.60e-09

The magnitude of the discrepancies were the same for CCSD and CCSD(T), indicating that the problem was unrelated to the MO symmetry issues the AI confidently hypothesized earlier.

CASE 1: CCSD(T) POLARIZABILITIES

AN ELEGANT SYMMETRY EXPLANATION THAT WAS SIMPLY WRONG

My experience suggested that this was simply a discrepancy in the set up of the calculation between the two codes:

- Inconsistent precision of the specification of the geometry?
- Differences in the convergence criteria for the wave function components?
- Discrepancies in the definitions of the basis set exponents and/or contraction coefficients?

I asked Claude to (a) give me the nuclear repulsion energies to at least 12 decimal places from each code; (b) to double-check that the convergence settings were the same between the two; and (c) to carefully compare the Psi4 and CFOUR basis sets.

Its finding: Psi4's 6-31G and CFOUR's GENBAS ("SOSA 6-31G FROM GAUSSIAN-82") differ in the sixth decimal of one hydrogen exponent, **2.8253944 vs. 2.8253937**, with the O and F coefficients stored to fewer digits.

CASE 1: CCSD(T) POLARIZABILITIES

AN ELEGANT SYMMETRY EXPLANATION THAT WAS SIMPLY WRONG

After adjusting Psi4's 6-31G basis set definition (from the Basis Set Exchange) to match the older definition in CFOUR:

component	pycc	CFOUR	|pycc – CFOUR|
α_xx	-2.2106648649	-2.2106648649	<1e-10
α_xy	-0.7220673958	-0.7220673958	<1e-10
α_yy	0.1661531582	0.1661531583	1.00e-10
α_zz	0.0652603834	0.0652603834	<1e-10

The two codes matched to nearly every decimal place printed by CFOUR in its output files.

The PyCC code was always correct.

CASE 1: CCSD(T) POLARIZABILITIES

AN ELEGANT SYMMETRY EXPLANATION THAT WAS SIMPLY WRONG

I asked Claude to give me its perspective on the take-aways from this episode, and I quote:

- **Plausibility is not evidence.** The σ_h story was coherent, quantitative and literature-anchored. **Its only defect was being false.**
- **The AI closes issues too early.** Twice it wrote some version of "it is not a bug" and "the code is correct". Both times the human declined.
- **"Good enough" agreement hides real problems.** Accepting 10^{-7} would have left an undiagnosed basis mismatch in the validation record. The last three decimal places are where the actual defect lived.
- **The human found the root cause.** The AI compared the files, but the hypothesis that the basis tabulations differed came from the PI.

CASE 2: CCSD(T) GRADIENTS

WHICH FORMULATION IS CORRECT?

The analytic CCSD(T) gradient involves the MO-relaxed one-electron density, the unrelaxed two-electron density, and the orbital rotation gradient:

$$\frac{\partial E}{\partial x} = \sum_{pq} \tilde{D}_{pq} f_{pq}^{(x)} + \sum_{pqrs} \Gamma_{pqrs} \langle pq \| rs \rangle^{(x)} + \sum_{pq} I_{pq} S_{pq}^{(x)}$$

The question to the AI: Do CCSD(T) gradients require canonical or non-canonical perturbed MOs?

$$\frac{\partial f_{pq}}{\partial x} = \begin{cases} 0 & \text{if } p \neq q \text{ (canonical perturbed-MOs)} \\ \frac{\partial f_{pp}}{\partial x} & \text{if } p = q \text{ (non-canonical perturbed-MOs)} \end{cases}$$

Answer: Both will work, but one yields greater computational efficiency.

CASE 2: CCSD(T) GRADIENTS

WHICH FORMULATION IS CORRECT?

In the effort to extend PyCC to include analytic CCSD(T) gradients, I engaged Claude in a discussion of potential formulations by providing with the two primary sources by Scuseria and by Lee and Rendell (both 1991), as well as a more recent formulation by Hald, Halkier, Jørgensen, Coriani, Hättig, and Helgaker ("the Vikings"), and a paper by Handy from the mid-1980s on MP2 gradients and Hessians.

In discussions of these papers and others over a couple of days, the AI asserted twice that non-canonical perturbed MOs were definitely **required**. Claude's central argument was that the CCSD(T) energy was **not invariant** to unitary transformations among the occupied or the virtual orbitals, and thus the non-canonical perturbed-MO formulation was the only possibility for correct results.

I'm embarrassed to admit that its arguments at one point almost had me convinced, even though part of my postdoctoral research involved exactly this issue of invariance!!

CASE 2: CCSD(T) GRADIENTS

WHICH FORMULATION IS CORRECT?

Analytic evaluation of energy gradients for the singles and doubles coupled cluster method including perturbative triple excitations: Theory and applications to FOOF and Cr_2

Gustavo E. Scuseria
Department of Chemistry and Rice Quantum Institute, Rice University, Houston, Texas 77251-1892

$$Q_{ij} = -\tfrac{1}{2}\mathcal{P}_{ij}M_{ij},$$
$$Q_{ab} = \tfrac{1}{2}\mathcal{P}_{ab}N_{ab},$$

$$M_{ij} = \sum_{klabc} (W+V)^{abc}_{ikl} R\,[W]^{abc}_{jkl},$$
$$N_{ab} = \sum_{ijkcd} (W+V)^{acd}_{ijk} R\,[W]^{bcd}_{ijk},$$

Scuseria, *J. Chem. Phys.* **94**, 442-447 (1991):

The (T) one-electron density contribution is **non-diagonal** in the occupied- and virtual-MO blocks.

Analytic gradients for coupled-cluster energies that include noniterative connected triple excitations: Application to *cis*- and *trans*-HONO

Timothy J. Lee and Alistair P. Rendell[a)]
NASA Ames Research Center, Moffett Field, California, 94035

$$Q^{\mathrm{CCSD(T)}}_{ij} = Q^{\mathrm{CCSD}}_{ij} + \epsilon_i \delta_{ij},$$
$$Q^{\mathrm{CCSD(T)}}_{ab} = Q^{\mathrm{CCSD}}_{ab} + \epsilon_a \delta_{ab},$$
$$Q^{\mathrm{CCSD(T)}}_{ia} = Q^{\mathrm{CCSD}}_{ia},$$

$$\epsilon_i = \sum_{jkabc} \left[2E^{abc}_{ijk} + E^{abc}_{jik}\right],$$
$$\epsilon_a = \sum_{ijkbc} \left[2E^{abc}_{ijk} + E^{bac}_{ijk}\right],$$

Lee and Rendell, *J. Chem. Phys.* **94**, 6229-6236 (1991):

The (T) one-electron density contribution is **diagonal** in the occupied- and virtual-MO blocks.

CASE 2: CCSD(T) GRADIENTS

WHICH FORMULATION IS CORRECT?

The answer to the question of canonical vs. non-canonical perturbed MOs is: **BOTH**. Both the Scuseria and Lee/Rendell formulations are valid, but the latter allows one to **eliminate** one of three extremely costly N^7 steps.

After re-convincing myself of the solution, I explained this to Claude in detail and instructed it on how to adjust our existing code to (a) use the Lee/Rendell **DIAGONAL** occ-occ/vir-vir onepdm **AND** (b) to include an occ-vir contribution (which allows correct calculation of unrelaxed dipole moments.

quantity	before vs. after
`E(T)`	4.3e-19
diag `Doo`	3.3e-19
diag `Dvv`	8.1e-20
gradient	4.9e-17
`t3_density` wall time	**Reduced by a factor of 3.25x**

CASE 2: CCSD(T) GRADIENTS

WHICH FORMULATION IS CORRECT?

Again, I asked Claude to give me its perspective on the take-aways from this episode, and I quote:

- **The theory came from the human.** The AI did not derive the canonical/non-canonical resolution; it had asserted the opposite twice and was corrected from the literature.
- **The AI was extremely good at the next step.** Given the correct physics stated clearly in prose, it located the exact redundant loop by line number, produced a correct code replacement, proved the permutational-symmetry argument for why the adjustment is correct, and validated the result to 10^{-17} with a measured 3.25x speedup. It did all of this in about 20 minutes.
- **Bottom line:** The AI was a poor theoretician and an outstanding research assistant.

Important side note: I realized after this adventure that my own "production level/optimized" CCSD(T) gradient code in Psi4 (from around 2005) used canonical perturbed-MOs, but still included the extra N^7 step! I'm in the process of removing the unnecessary computational expense.

SOME PERSONAL OPINIONS ON THE USE OF AI IN QUANTUM CHEMICAL RESEARCH

- **Good Use of AI:** To carry out complicated tasks that you know very well how to do, but that the AI can do faster and do not inherently require a human's effort.

- **Bad Use of AI:** To carry out complicated tasks that you don't really understand, but that you want to get done quickly and hope that no one will find out that you didn't do it yourself.

- **Extremely Good Use of AI:** To carry out a complicated task that you understand somewhat, but would like/need to understand much better. The challenge: recognizing when you genuinely do and do not understand – and also recognizing when the AI is giving you incorrect information.

SOME PERSONAL OPINIONS ON THE USE OF AI IN QUANTUM CHEMICAL RESEARCH

- **Do not let the AI agent get away with responding to your questions above your understanding.** AI agents (especially Claude) tend to give high-level, verbose responses with lots of jargon and technical detail. Do not let them do this. Tell the agent your background and require it to explain it to you at your level of knowledge.

- **Do not immediately trust the AI agent's answer.** It can be tempting to simply accept its response at face value (because you're in a hurry, or you're unwilling to admit to yourself that you don't quite get it). AI agents can make mistakes because they ultimately rely on human knowledge and/or they are willing to extrapolate beyond the given evidence.

- **When you might disagree with the AI, find a relevant check for which you know the answer.** Before trusting it on what you do not know, ask it about something you know cold: a paper you wrote, a derivation you have done by hand, a result you can look up.

- **The AI's agreement with you is not evidence that it's right.** Push back on a correct answer and it will frequently abandon it because agreement and praise are the cheapest things it produces.

SOME PERSONAL OPINIONS ON THE USE OF AI IN QUANTUM CHEMICAL RESEARCH

- **Verify every reference.** AI agents will sometimes produce citations that look right and do not exist, or that exist but do not say what they claim.
- **Never give the AI what is not yours to share.** Unpublished data, a manuscript you are refereeing, a collaborator's draft, proprietary code. This is a professional-ethics obligation independent of whether the AI tool is any good.
- **You are responsible for everything that carries your name.** If you have to say “the AI produced it” as a defense, then you’re not really defending yourself.
- **Writing is thinking, so outsourcing the writing outsources the thinking.** If the AI has to write it for you, then you do not understand what you’re trying to write.
- **Do not let it choose your research direction.** AI agents tend to be good at giving conventional answers, but not very good at giving original ones.