NU 520: Quantum field theory for mathematicians (Winter 2026)
Syllabus

Course Description: A (rigorous!) introduction to quantum field theory targeted towards mathematicians. The first half of the course will be devoted to the full construction of free fields, including non-scalar fields and massless fields. For this part of the course, we will follow Weinberg's *Quantum Theory of Fields*, Volume I, Chapters 2–5 and Chp. 9–10 of Talagrand's *What is a Quantum Field Theory?*, supplemented by my own notes.

The second half of the course will be an introduction to (perturbative) interacting theories, focusing on quantum electrodynamics (QED) as the paradigmatic example. QED is often touted as the most precisely verified theory in science. The QED prediction for the anomalous magnetic moment of the electron matches experiment to more than ten (!) significant figures. The goal of this course is to get to the algorithm via which these predictions are derived, *without ever writing down an ill-defined integral.* Emphasis will be placed on the physical principles from which that algorithm is derived.

By this point, the standard textbook treatment has abandoned rigor. There is one exception among textbooks following this line — Scharf's *Finite QED* (Chp. 3 & 4) — on which our treatment will be based, except I'll translate Scharf's computations into more familiar terms using Feynman diagrams. The formalism of Bogolyubov–Stückelberg, often known as "causal perturbation theory," is used to axiomatize the scattering matrix (S-matrix). This makes things rigorous, but only at the level of perturbation theory. The existence of solutions of the axioms, roughly amounting to the solution of the renormalization problem order-by-order in perturbation theory in the traditional approach, will be proven by microlocal methods, which will be introduced as needed. Microlocal analysis is *not* a prerequisite for this course.

Topics: A tentative, rather optimistic, list of topics is as follows:

1. Projective representation theory of the Poincaré group. (Weinberg Chp. 2)
2. Wave mechanics and wave-mechanical realizations of the Poincaré irreps. (Weinberg Chp. 5, own notes)
3. The "second quantization" functor. Free quantum fields (inc. the spin-statistics and CPT theorems, for free fields). (Weinberg Chp. 5)
4. Coherent states of bosonic fields and the emergence of the macroscopic electromagnetic field. (Own notes.)
5. The scattering matrix (S-matrix); cross-sections, decay rates, etc. The cluster decomposition principle. (Weinberg Chp. 3, 4)
6. The axiomatization of the S-matrix in QFT: the Bogolyubov–Stückelberg axioms. (Scharf Chp. 3, own notes.)
7. QED to one-loop. (Scharf Chp. 3, own notes.)
8. The Epstein–Glaser existence theorem; the existence of solutions of the axioms. QED to all orders. (Scharf Chp. 4.)

I'm imagining spending slightly more than one week on each topic, on average. Time willing, we can discuss the difficulties in going beyond perturbation theory:

9. The divergence of the perturbation series, Borel resummability, renormalons, and instantons. The Yang–Mills millennium problem.

What this course is not: The guiding philosophy of this course is to make a beeline to the QED scattering matrix. Consequently, there are many topics that are often covered in similar courses that we will eschew completely. This includes the canonical formalism and equal time commutation relations, anything based on Euclidean QFT or path integrals, effective field theory, the renormalization group, the net of local observables, the Wightman axioms and functions, Green/time-ordered functions and the Gell–Mann and Low theorem, Haag–Ruelle scattering theory and LSZ reduction, or rigorous constructions of QFTs in low-dimensions. More complicated gauge theories, such as Yang–Mills/QCD/the standard model will not be broached. One cannot cover all of these things in a single quarter anyways. However, it is hoped that this course can provide a point of departure for those interested in these diverse topics.

Prerequisites: We will make use of a smattering of topics from a standard mathematics graduate curriculum, but there are no formal prerequisites. I will make heavy use of the language of distributions, so it will be good to know the rudiments of that theory.

As far as physics is concerned, there are no prerequisites, though previous exposure to quantum mechanics and electrodynamics will undoubtedly be helpful.

Teaching Staff:

Instructor: Ethan Sussman
E-mail: `ethan.sussman@northwestern.edu`
Website: `ethansussman.github.io`

Lecture: Mon, Wed, Fri 3:00PM - 3:50PM

Course website: To be hosted on my website.

Grading: Students enrolled for a letter grade will be asked to complete a few short problem sets, which will be graded on completion.

FAQ.

- **Q. What about renormalization? A.** The infamous problems of *renormalization* never rear their head in the Epstein–Glaser approach. Instead, the leading-order term (the interaction Lagrangian) only determines higher-order terms as functionals of the volume cutoff away from the diagonal. So, instead of having an ill-defined theory, one has a well-defined mathematical problem about extending the distributions in the position-space Feynman rules to various subvarieties. There exist multiple extensions. The difference between any two extensions is a "renormalization ambiguity." These form a finite-dimensional space of distributions.

  The problem of performing the extension was rigorously solved by Epstein–Glaser. Rather than following the original Epstein–Glaser approach (as Scharf does), we will deploy more modern microlocal tools, introduced into this subject by Brunetti & Fredenhagen in order to do QFT on curved spactimes. This will allow us to phrase the constructions in terms of position-space Feynman diagrams. Unlike Epstein–Glaser's presentation, this comes quite close to physicists' presentation, but automatically includes regularizers making the position-space Feynman integrals well-defined. The necessary microlocal tools will be developed or blackboxed as needed.

- **Q. Doesn't the S-matrix have IR divergences that render it inapplicable as a foundation? A.** The Bogolyubov–Stückelberg S-matrix has a built-in volume (a.k.a. IR) cutoff, and what is axiomatized is how the operator depends on the cutoff.* It therefore requires a solution to the "ultraviolet" (UV) problem as part of the definition of the theory, but not a solution to the "infrared" (IR) problem. The IR problem can now be phrased precisely about the limit of various well-defined quantities as the IR cutoff is removed.

  Rather than the S-matrix, the cynosure of QFT is sometimes the net of observables of the interacting fields and sometimes the correlation (Green or Wightman) functions of the interacting fields. A formula of Bogolyubov allows the interacting fields (without an IR cutoff!) to be recovered from the S-matrix (with external fields, but still with an IR-cutoff). Correlation functions follow, except the IR cutoff is still present. However, unlike the S-matrix, the correlation functions converge as the IR cutoff is removed, as proven by Epstein–Glaser (for massive fields), Blanchard–Seneor (for QED), and Duch (in general). (Remember that this is all at the level of perturbation theory.)

  So, as far as constructing the theory goes, it suffices to define the S-matrix *allowing an IR cutoff*. We will not talk about correlation functions or the net of observables in this course.

  The IR problem for the S-matrix is only partially solved (and remember that this is only at the level of perturbation theory), but this is not special to causal perturbation theory. In conclusion, we will not worry about IR divergences in the S-matrix, and this is not a lapse of rigor.

- **Q. Don't you need path-integrals do have a manifestly Lorentz-invariant perturbation theory? A.** Besides rigor, causal perturbation theory has another advantage: it is manifestly unitary *and* Lorentz invariant. In his textbook, Weinberg makes much of certain non-covariant "seagull" terms that arise at second-order in Dyson's formula in theories with spin ≥ 1 fields. Weinberg lands on a dual formalism: the canonical formalism to make unitarity manifest, and the path-integral formalism to make Lorentz-invariance manifest. However, in the Epstein–Glaser approach, the problem posed by seagull terms is dealt with simultaneously with the solution of the UV problem: extending a Lorentz-covariant distribution on $\mathbb{R}^{4N}\backslash\{\text{origin}\}$ to a Lorentz-covariant distribution on $\mathbb{R}^{4N}$. One can show on general grounds that if an extension exists at all, it is always possible to choose a Lorentz-covariant one. Something similar applies to gauge symmetry in gauge theories (see Scharf's *A True Ghost Story* for the non-Abelian case).

*This is morally similar to functorial QFT, in which a perturbative QFT is a certain functor from a category of cobordisms with appropriate extra structure. The Bogolyubov–Stückelberg formalism is much simpler, technically speaking.