Gauge invariance applies to statistical mechanics too

Mathematical tools from the abstract world of quantum fields have surprising relevance to the seemingly more concrete realm of particles in boxes.

To the uninitiated, the standard model of particle physics can seem like a random hodgepodge of particles and forces: quarks and gluons, charged leptons and neutrinos, W and Z bosons, each with its own idiosyncratic properties and behaviors. To paraphrase I. I. Rabi's remark about the muon, "Who ordered that … or that … or that?"

But a deeper dive into the theory reveals a method to the madness. Far from being arbitrary, many features of the model follow mathematically from the symmetries of the universe. Once the symmetries are known, much of the rest follows inevitably.

The theoretical workhorse for deriving physical laws from symmetries is the gauge transformation. Roughly speaking, you start with a quantity, such as the phase of a quantum mechanical wavefunction, that doesn't affect any physical observables, and you write it as a local function that takes different values at different points in space. Turn the mathematical crank, and out pops a physical law—in this case, a description of the existence and behavior of photons.

Now, Matthias Schmidt and colleagues at the University of Bayreuth in Germany have shown that gauge transformations can also be fruitful in a seemingly disparate area of physics: statistical mechanics.[1] They're still exploring all the consequences of their discovery, but they've already uncovered a plethora of mathematical structure, equations that can help to characterize soft-matter systems, and questions about what statistical mechanical averages really mean.

Mindset shift

It all started with an offhand remark in 2019. Schmidt was working with Sophie Hermann, a new PhD student in his group, to explore the effect of a mathematical manipulation that he called "shifting." "Sophie is a very clear and

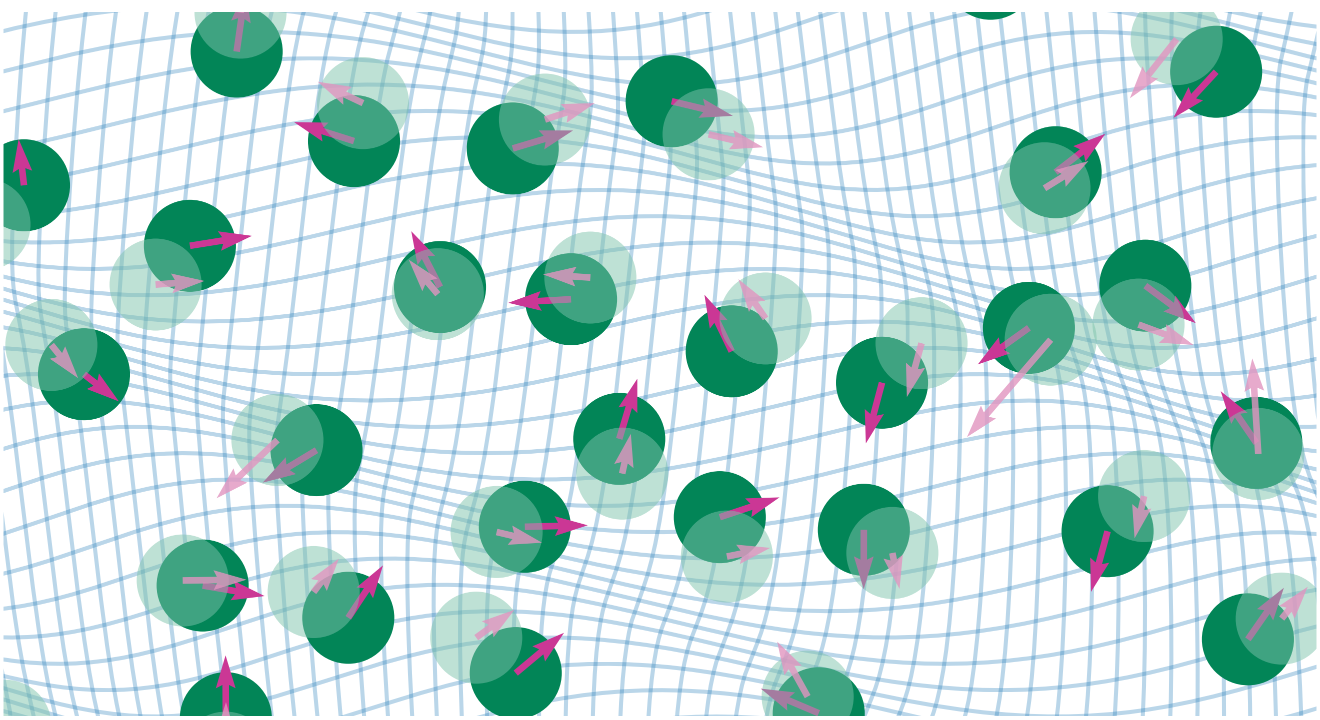

FIGURE 1. A SHIFTY TRANSFORMATION. In a statistical mechanical ensemble at equilibrium, when the position **r** of each particle (solid circles) is shifted to a new position (transparent circles) by a smooth vector field $\boldsymbol{\varepsilon}(\mathbf{r})$ and the momenta are adjusted in a corresponding way (solid and transparent arrows), the values of all observable quantities remain unchanged. The realization that the shift is a symmetry of the system—a gauge transformation—can be used to derive new equations about how the particles behave. (Figure courtesy of Florian Sammüller.)

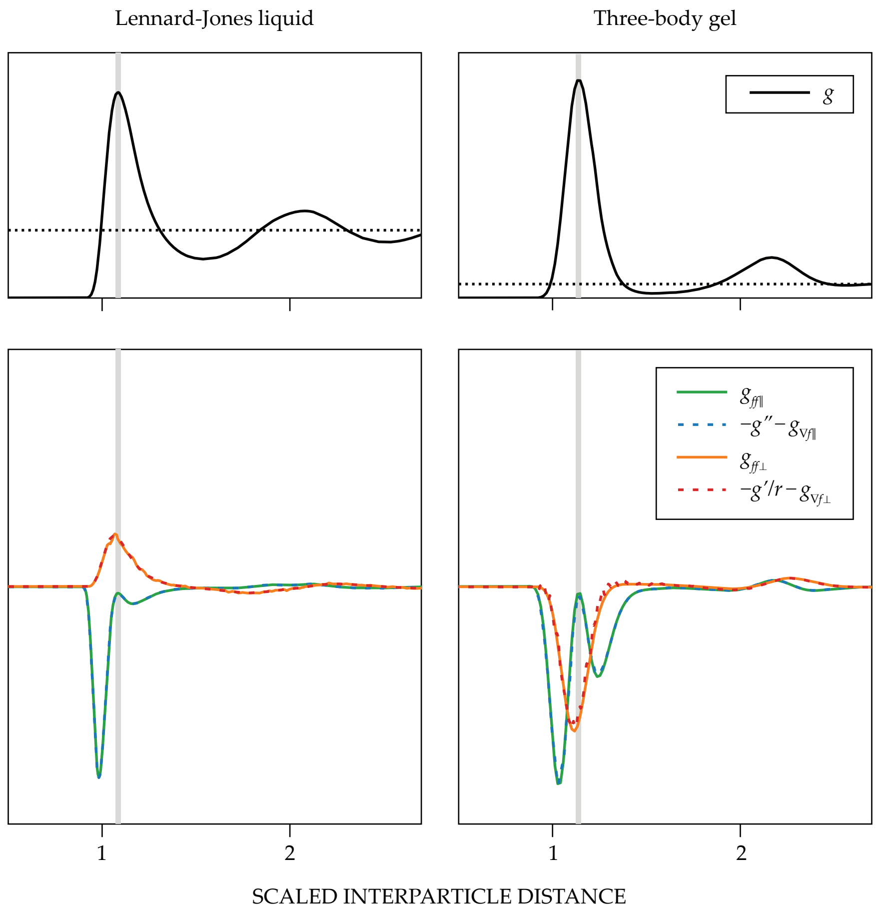

FIGURE 2. LIQUID OR GEL—HOW TO TELL? The position correlation function g can be qualitatively similar for different forms of soft matter, including common models of a liquid and a gel, as shown by the black curves in the upper two panels. But the gauge transformation from figure 1 generates equations involving other correlations among forces (g_{ff}) and force gradients ($g_{\nabla f}$) that can be more sensitive to a system's macroscale mechanical properties. Gauge invariance implies that the quantities shown by the solid and dashed lines in the lower panels should be equal; the plots, derived from numerical simulations, show that the theoretical predictions are correct. (Adapted from ref. 3.)

systematic thinker," says Schmidt, "and she kept insisting that it was unclear what 'shifting' actually implies."

Grasping for an answer, he appealed to a topic he'd covered in his undergraduate classical mechanics course, which Hermann had taken a few years previously: "Think of it like using Noether's theorem," he said. "Translational invariance in a given direction implies conservation of momentum in that direction."

Schmidt was referring to Emmy Noether, the foremother of modern thinking about the role of symmetry in physics. With her theorem, published in 1918, she proved that whenever a system is invariant under a continuous symmetry, it has a corresponding conserved quantity. Translational symmetry implies conservation of momentum, rotational symmetry implies conservation of angular momentum, and time-translation symmetry implies conservation of energy.

Those undergraduate-friendly examples might seem pedestrian and hardly worth mentioning, but the theorem's implications go far deeper. Noether herself was drawn to the problem by the desire to reconcile what physicists thought they knew about classical mechanics with the new theories of special and general relativity. A relativistic universe—especially if it's expanding—might not be translationally or time-translationally invariant, so it might not conserve momentum and energy. But it has other symmetries, and thus other conserved quantities. Noether laid the foundations for understanding it all.

"It was meant to be a throwaway comment," says Schmidt. "What I hadn't expected was that Sophie would go back to Noether's original paper, work through it, and come back with the conclusion that the idea actually has some real substance in it. Once that was clear, we just sat down and worked it out as clearly as we could."

To start with a simple example, they considered shifting the position $\mathbf{r}$ of each particle in an ensemble by a constant vector $\boldsymbol{\varepsilon}$. That's not inherently a symmetry of the underlying classical mechanical system, because they envisioned the particles moving in an external energy potential $V(\mathbf{r})$ that stays put under the shift. So when the particles' positions change, their energies do too. But when the researchers looked at the system as a statistical mechanical ensemble, something more subtle happened.[2]

The basic operation of equilibrium statistical mechanics is the computation of weighted averages by integrating over all possible arrangements of individual particles, with each arrangement, or "microstate," weighted by $e^{-E/kT}$, in which E is the total energy, T is the temperature, and k is Boltzmann's constant. The weighting reflects the fact that low-energy microstates always show up with the highest probability, but the higher-energy ones are not ruled out, especially at higher temperatures.

Shifting a microstate changes its energy and, therefore, its weight in the average. As a result, it turns out, the equilibrium average—of any observable quantity—is unaffected by the shift. Shifting by $\boldsymbol{\varepsilon}$ is not a symmetry under classical mechanics, but under statistical mechanics, it is.

Mathematically, the symmetry means that in thermal equilibrium, the derivative $dX/d\boldsymbol{\varepsilon} = 0$, no matter what X is. Hermann and Schmidt took X to be $\sum V$, the sum of the total potential energy of all particles—and the derivative of potential energy is just the force exerted by that potential. Ergo, in equilibrium, $\sum \mathbf{F}_{ext}$, the sum of external forces on the system, equals zero.

That might seem boringly obvious. If $\sum \mathbf{F}_{ext}$ were not zero, the system would start to move, which would mean it hadn't been in equilibrium after all. But as the researchers pointed out, $\sum \mathbf{F}_{ext} = 0$ is not true for most of the individual microstates. Rather, it's a nontrivial statement about the nature of thermal equilibrium—a so-called sum rule—and Noether's theorem offered a new way of proving it.

From global to local

With subsequent waves of group members over the past five years—including Florian Sammüller and Johanna Müller—Schmidt and Hermann continued to develop the theory. In particular, says Schmidt, "Noether's theorem comes in two flavors, local and global. We'd started with global shifts, but the real powerhouse is the local version."

Generalizing from global to local symmetry would mean changing $\boldsymbol{\varepsilon}$ from a single vector to a position-dependent function $\boldsymbol{\varepsilon}(\mathbf{r})$. In general, shifting by $\boldsymbol{\varepsilon}(\mathbf{r})$ is not a statistical mechanical symmetry: The shift spreads out some particles and moves others closer together. The distortion leaves the respective microstates either overrepresented or underrepresented in the integral over all microstates.

But the thermal average is an integral over not just the particles' positions but also their momenta. So the Bayreuth researchers introduced a corresponding momentum transform, as shown in figure 1, that compensated for the effect of the position shift: Where $\boldsymbol{\varepsilon}(\mathbf{r})$ spread the positions apart, the momenta were correspondingly compressed, and vice versa. As a result, the position–momentum shift (still referred to as simply $\boldsymbol{\varepsilon}(\mathbf{r})$ for brevity) once again left the equilibrium averages of all observables unchanged.

Unraveling the consequences of the local symmetry follows similar lines. Now, $dX/d\boldsymbol{\varepsilon}(\mathbf{r})$ is what's called a functional derivative—a derivative with respect to a function—but just like an ordinary derivative with respect to a number, it can still be set to zero for any observable X. Moreover, one can study the second derivatives with respect to $\boldsymbol{\varepsilon}$ to generate higher-order sum rules that relate the spatial correlations among forces and other quantities.

"All these sum rules just say, 'Zero equals zero,' " says Schmidt, "or 'These two things add up to zero,' where one is an obvious everyday object, and the other is some strange correlation function that you'd never otherwise think of measuring. But it's really worth it to study them, because they can tell you a lot about the system you're looking at."

For example, for their first foray into exploring the consequences of the local symmetry, they looked at simulations of liquids and gels.[3] Those two forms of matter have obvious differences on the macroscale, but it can be tricky to relate their properties to what's happening on the microscale. "The natural thing to want to measure is the shell structure—how likely particles are to be some distance apart," says Sammüller. That quantity, plotted as g in the upper panels of figure 2, is qualitatively similar between model liquids and model gels.

But when the researchers differentiated energy twice with respect to $\boldsymbol{\varepsilon}(\mathbf{r})$, they got a sum rule relating derivatives of g to correlations of forces g_{ff} and force gradients $g_{\nabla f}$. The correlations would be hard to measure in real fluids, but they're certainly measurable in simulations and possibly even in experiments on micron-sized colloids. And as the bottom panels show, they're starkly different between liquids and gels, and the quantities that the sum rule predicts to be equal really are.

"These quantities that come out of the analysis can be very sensitive to various important physical mechanisms," says Sammüller. "They might even be useful for designing liquids with tailored properties."

Full circle

"We could have continued like this," says Schmidt, "with a new paper for every observable: 'Now we can do this for energy, now for kinetic energy,' and so on." But when Müller joined the group, she brought with her a master's degree in mathematics—and the tools to show just how general the shifting theory really was.[1]

The universe of all possible shifts $\boldsymbol{\varepsilon}(\mathbf{r})$, it turned out, forms a mathematical structure called a Lie algebra (named after Norwegian mathematician Sophus Lie—nothing to do with prevarication). Lie algebras turn up in many other areas of physics and mathematics, including in the gauge transformations from particle physics. "Dealing with gauge invariance and Lie algebras is such a standard thing in other areas," says Schmidt, "and it helps us to better understand, assess, and manage the implications of the mathematics."

In particular, the Lie algebra structure sets clear boundaries on the types of sum rules that the $\boldsymbol{\varepsilon}(\mathbf{r})$ shifts can generate. No matter what observable quantity the researchers start with or which functional derivatives they calculate, they'll end up with a sum rule involving correlations of forces and other specific force-like quantities. "These do form a hierarchy of increasing complexity, but the complexity is within the limits set by the Lie algebra," says Schmidt. "The sum rules don't proliferate into an uncontrolled, ever-increasing range of quantities that they relate to each other."

The implications of the $\boldsymbol{\varepsilon}(\mathbf{r})$ shifts were falling into place, but there remained the matter of Hermann's original question: What does the shift really mean? "Gauge invariance is a brutal thing somehow," says Schmidt, "because it says that all these things that one can reach with the gauge transformation are really the same." That is, the gauge transformation is more than a mere mathematical manipulation: The transformed and untransformed versions of the system are physically indistinguishable, which means they're also physically equivalent.

Other common targets of gauge transformations, such as quantum fields and electromagnetic potentials, are already such abstract entities that it's relatively easy to accept that one way of writing them down is no more physically real than any other. Statistical mechanics seems different in that regard, because classical intuition gives rise to mental movies of ensembles of particles zipping around in boxes. Those microstates might seem too concrete to exist as part of an $\boldsymbol{\varepsilon}(\mathbf{r})$-shifted equivalence class: Surely one set of positions and momenta must be the real one?

"It's absolutely weird that gauge invariance also applies in this context, and it's hard to get your head around," says Schmidt. "But it's the averages we're taking that are the abstract thing. The movies aren't real—they're just one very specific illustration. It's possible to look at a system too accurately."

Johanna L. Miller

References

1. J. Müller et al., *Phys. Rev. Lett.* **133**, 217101 (2024).
2. S. Hermann, M. Schmidt, *Commun. Phys.* **4**, 176 (2021).
3. F. Sammüller et al., *Phys. Rev. Lett.* **130**, 268203 (2023).