Report

Virasoro TQFT is a machinery that can compute the partition function of three-dimensional quantum gravity on manifolds of fixed topology. It is much more convenient that a direction computation in a metric formulation, since it can be done for arbitrary hyperbolic three-manifolds while the metric formulation quickly runs into limitations. Perhaps even more interestingly, it can be done at finite central charge. This means it repackages arbitrary loop corrections. The authors build on their earlier work which they briefly summarize and discuss applications to holography: the partition functions computed on a given topology have an interpretation in terms of a boundary dual given by an ensemble of CFT data (a spectrum and OPE coefficients). In particular, the bulk partition functions determine moments of the OPE coefficients. The authors go through multiple examples, discuss the bulk computations and their associated boundary interpretation. They focus on the non-Gaussian corrections to the distribution of OPE coefficients, and show how they are required for a crossing-symmetric boundary ensemble. Finally, they discuss the figure-eight knot complement, whose partition function they evaluate using the Virasoro TQFT (in multiple ways) and compare it to the calculation in Teichmuller TQFT, finding agreement.

This is well-written paper with a clear message, and the calculations appear to be correct to me. The authors go through multiple examples, and also make connections with pre-existing results in the literature, giving additional support for their proposal. I do have some small comments that I would like the authors to address before I recommend the paper for publication

- Below equation (2.3), the expression $\Sigma_{0,3}$ has not been defined.
- I found both interesting and confusing the concept of mutations, and the fact that the Virasoro TQFT partition functions cannot completely distinguish a manifold. What worried me, is a "double-counting" type of problem. Are there cases where given a fixed boundary data, we have two bulk manifolds (or more) which are distinct but have the same partition function? Can this not lead to an expression for the full partition function (i.e. the sum over topologies) where the vacuum is counted with multiplicity greater than one?
- The authors discuss various examples, some of which correspond to the same boundary condition but where two different bulk topologies contribute (and in principle compete). For example, there is (3.55) and (3.61). Is there no way to guess which geometry dominates? A naive guess would be: the one with the smallest phase, since the phase will induce a high oscillation making integrals small. Is this intuition wrong? Are there counterexamples?
- Finally, what is the holographic interpretation of the figure-eight knot complement? The authors do not seem to discuss this, and it seems like a natural question to ask.