Regarding the submission 10836 to NeurIPS 2020, we are grateful to the reviewers for their time and efforts for understanding our work, appreciating the strengths, expressing their concerns and enlightening us with valuable suggestions. We further appreciate NeurIPS for allowing us to clarify on the concerns raised.

Reviewer 2: Computer vision results not explicitly shown, also code-data release is missing: We second this concern that any discrepancy in capturing the input from the environment will impact the performance. Since the functionality of YOLO was customized towards developing region traffic and evaluated in prior work (reference [3] in the paper), we verified our traffic density outputs using the results of [3] as ground truth. Also, code will be released for all modules, along with relevant data, upon acceptance as committed in the paper.

Reviewer 2: Additional results from developing countries: The extended results, that were held down due to space limitation, are given in Table 1. We presented the results of most heavy traffic data in the paper. We will try to include 1-2 more experiments, or at least mention that we have tested at the same intersection (due to logistic constraints of camera deployment), but at other times under different traffic conditions, and observed similar improvements.

Table 1: **Performance of EcoLight Algorithms on 1x1 Traffic intersection (at different times)**

	2			3			4			5		
Algo	nOut	Travel	Total	nOut	Travel	Total	nOut	Travel	Total	nOut	Travel	Total
PressLight	1300	222.0	308.7	225	29.9	29.8	529	163.1	177.6	557	161.8	189.4
Colight	1394	207.9	280.5	221	33.1	51.2	514	163.2	181.7	514	159.0	208.2
2dimRL	1418	276.3	277.8	225	30.8	30.4	540	182.3	178.6	574	182.4	189.0
FairShare(T)	1417	266.9	277.6	225	31.0	30.4	517	194.5	189.5	581	193.5	188.3
Decision(T)	1466	273.7	264.2	225	31.0	30.6	538	185.2	180.2	579	176.6	186.5
FairShare(TS)	1421	289.1	278.0	224	30.6	30.1	538	181.7	178.4	587	191.8	187.4
Decision(TS)	1487	266.1	256.5	225	30.6	30.0	564	172.8	167.8	584	158.7	183.2
Timed(2dim)	1496	266.6	257.3	225	31.7	31.2	566	168.9	162.5	603	179.7	177.8
Timed(1dim)	1476	268.6	258.8	225	31.7	31.2	563	174.8	166.9	598	184.5	181.5
Random	1476	270.0	260.4	225	31.7	31.2	562	172.8	166.3	559	183.5	182.5
SOTL	1462	279.4	269.4	223	50.8	53.9	539	196.0	186.7	580	199.2	192.2
MaxPressure	1248	334.0	319.8	224	46.6	45.6	487	205.5	202.0	509	225.6	216.7
FixedTiming	1380	302.3	289.1	224	67.9	66.1	512	198.4	188.8	550	202.7	194.0

Reviewer 3: Formal definitions absent: We intended to start our paper by defining our problem as MDP, as the reviewer suggested. It would be similar to Section 3 of reference [14] in the paper. Due to space limitation, we eased on this section considering it a known formulation in the domain, favouring other sections for effective space utilization. We request the reviewer to consider a trimmed version of section 3 of [14] that we will include in our paper, as a fulfillment of the suggestion. We cannot express the equations in this document as well, due to space constraints.

Reviewer 3: Why do small states work? Is this generalizable to larger road networks? We observe that high dimensional states sometimes cause a qualified DNN to overfit. Carefully designed small states exploiting domain knowledge, retaining just enough required information, exhibit both better fitting and better generalization i.e. allow DNN to hold better understanding of the environment. We have analyzed this over a range of experiments, and the results are encouraging. Our work which carefully analyzes low dimensional states and what domain knowledge to use, is a research paper in its own capacity, and is awaiting review at a suitable conference. We avoided referring to that ARXIV-ed paper for anonymity reasons and will include a reference in this paper, if it is accepted. Hence, we started the given EcoLight paper with good confidence in low dimensional states. In terms of generalizability, all our methods are single intersection based, without any communication dependency, and thus should be generalizable to any number of intersections in large road networks, each intersection's agent working independently. This is reflected in our good results on competent open-source multi-intersection traffic data (16x1 and 16x3), also used in prior arts.

Reviewer 5: Generalize optimized control for applications other than traffic: For the LUT oriented work, the quantization was employed to facilitate efficient inference at runtime, despite of the repercussions it might cause to the performance, and our Goodness based RL methods effectively compensated for the same. These RL methods, independent of any quantization limitation, can also be utilized to improve fairness of any state-of-the-art RL/DNN design, exploiting the beauty of a sound heuristic. The enthusiastic reader-researchers would be able to deploy either/both of these independent features as per their requirements. We will also explore these other application domains, in context of resource constrained environments, as part of our future work. We will mention the same in the paper.

We acknowledge that due to limited space, we could not express everything we intended to. We had to sacrifice some background information in favour of experimental analysis and new contributions. If given an opportunity to restructure the paper (with/out increased space) after acceptance, we would love to incorporate your suggestions as best possible.