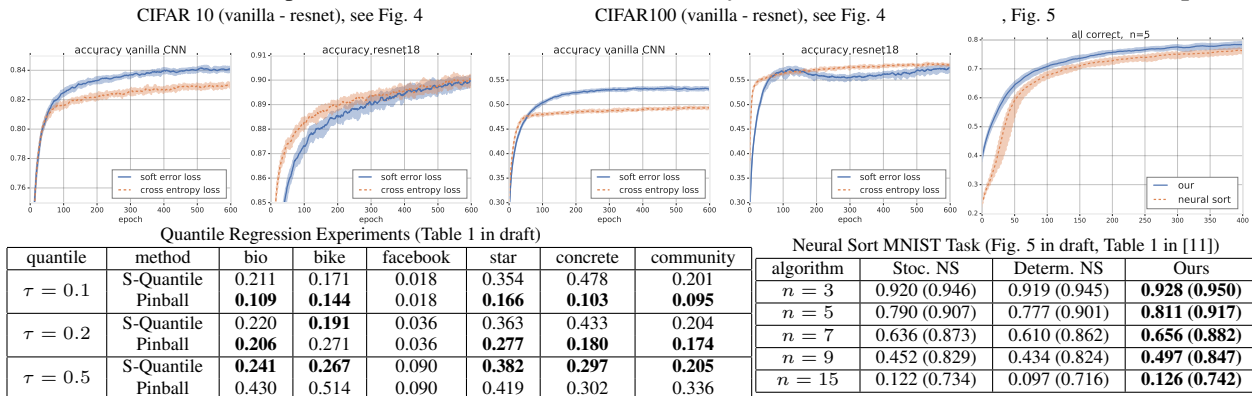


1 We thank all reviewers for their encouraging comments, and we apologize for many typos, notably in Def.2. Reviewers
 2 have all expressed similar reservations about experiments: We address these first and follow next with detailed answers.
 3 **Experiments:** We made an important finding when re-considering the default settings we set for the various parameters
 4 of our operators: $m, \mathbf{b}, \mathbf{y} \in \mathbb{O}_m$ (1.128), ε, ℓ (1.181), cost h (1.207) and rescaling of input values (1.214). Although
 5 all of these choices converge to standard sort (regardless of \mathbf{y}, h or of the rescaling) as soon as $n = m, \mathbf{a} = \mathbf{b}$ and
 6 $\varepsilon \rightarrow 0$, these choices do impact the “shape” of the differentiability encoded in our operators when $\varepsilon > 0$ (see p.6). Our
 7 finding is that our initial choice to do a **min-max** or a **softmax rescaling** of the input arrays into $[0, 1]$ as in 1.214 was
 8 **not a good choice**, leading to **instabilities and “squeezed” values**. Applying instead a **logistic map on standardised**
 9 **input values** (in batchnorm fashion) leads to **improved results across the board** (see below): Vanilla CNNs trained
 10 on CIFAR-10/100 with the soft-error loss (1.248) beat on average those trained with XE; RESNETs yield comparative
 11 results with both losses; In regression, using the soft-quantile loss (Eq.4) yields SOTA results when minimizing median
 12 error ($\tau = 0.5$), but mixed for small $\tau = 0.1, 0.2$ (current understanding: soft-quantiles are computed on mini-batches,
 13 therefore our loss is differentiable but *biased*, while the pinball loss is *not*-differentiable but *unbiased*). Finally, we beat
 14 neuralsort [11] on the setup shared in their colab. **We are now ready to share our code, for others to build upon.**



15 **Reviewer #1:** ▶ *applications did not seem overly original [...] were not the most natural.* We picked these tasks because
 16 they are representative of what can be done when switching to a rank/quantile based perspective. As envisioned by R2
 17 and R3, we also believe many original applications will follow. ▶ *not intimately familiar with works [1,11,15,16] [...] very little about [1,15,16]* We will clarify this difference. These papers use Sinkhorn to define a differentiable mechanism
 18 to produce a bistochastic matrix (as a relaxed permutation) from high-dimensional inputs, using a parameterized matrix-
 19 valued function (e.g. \mathbf{A} to $\mathcal{Z}^i(\mathbf{A})$ in p.6 of [1]). We use Sinkhorn + the maths of OT in 1D to define soft-CDFs and
 20 soft-quantiles **vector outputs**, as defined in Def.2. ▶ *the experimental results were [not] terribly supportive* We hope
 21 our new results are more convincing. ▶ *[...] helpful to give an example* We will work out an entire example in a
 22 blog post. ▶ *equations were typically far more complex than the simple ideas involved.* We respectfully disagree: the
 23 flexibility of our framework requires this level of detail, notably when $\mathbf{a} \neq \mathbf{b}$ (see for example 1.234). Note that these
 24 equations are exactly what we implemented in our code. We do however agree that the paper is currently hard to digest
 25 for casual Neurips readers. In order not to overwhelm them with technical details, we will move all of the OT related
 26 material from the intro in §1 to §2. Instead, we will introduce directly our work as “black-box” differentiable plug-ins
 27 for sorting, and encapsulate all the OT discussion in §2,3. ▶ *[...] top row (uniform, presumably?)* You are right.
 28 ▶ *Typo for \tilde{Q}^1 in Def. 2* You are right. A transpose T and a \circ are also missing before and after K . Apologies.
 29 **Reviewer #2:** ▶ *I think the basic idea is to do stuff like this [...].* Your pseudo-code agrees with the spirit of our work.
 30 However, we argue that our implementation is far more effective, owing to the flexibility given by weights \mathbf{b} and
 31 number of targets m . Writing $n = \text{len}(Xs)$ and $\ell = n_iterations$, your implementation requires $n^2 \cdot \ell$ operations.
 32 Ours has linear complexity $n \cdot m \cdot \ell$ (see 1.225). When computing a single quantile—the case you consider in your
 33 pseudo-code—we bring m down to 3 (see 1.234, and bottom row of Fig.3 where $m = 5$). Finally, we integrate directly
 34 quant within \mathbf{b} , no need to round to get an index (see Figs. 2 & 1(b)). ▶ *I actually found the motivation distracting.*
 35 As mentioned to R1, we have moved references to OT from §1 to §2,3. The rigorous derivation of our tools is now
 36 encapsulated in §2,3, which can be skipped on a first read. We will “sell” in §1 our operators as “no-brainer” plug-in
 37 replacements to regular sort. ▶ *\mathbb{O}_n seems really important [...] like to see this discussed.* Indeed, it is crucial that
 38 $\mathbf{y} \in \mathbb{O}_m$ to recover the soft-ranks/sorts of \mathbf{x} . This is discussed in 1.93-97 (we will expand). This is the “magic” from OT
 39 theory that makes everything work! (magic revealed in the proof of [Theo 2.9,23] ▶ *Results. Not great, but that’s ok.*
 40 We hope our new results are more convincing. ▶ *Some like algebra, some like code.* Agreed, we will add more code.
 41 **Reviewer #3:** ▶ *More convincing numerical experiments would certainly improve the paper.* We are excited to report
 42 more convincing experiments. ▶ *[...] considered certain algorithms that involve sorting as intermediate steps.* Your
 43 point illustrates perfectly our motivation and we are genuinely excited by the opportunities given by this new tool. For
 44 instance, quantile-based losses are related to fairness (e.g. arxiv 1907.08646) and we are keen to investigate connections.
 45