We deeply appreciate all the insightful review comments. We will fix all writing glitches, improve clarity and quality of writing, correct the confusions, and cite the missing references in our final paper with the major issues responded below:

Discussion of previous works: We briefly comment on some references below and discuss more in our revised paper.

[1] proposes a surrogate gradient BP method called Superspike. It uses the partial derivative of the negative half of a fast sigmoid as the surrogate gradient function to circumvent the non-differentiability of spikes. In addition, the authors also investigated different feedback methods to generate error signals from the output layer to hidden layers.

[2] presents a BP method for recurrent SNNs based on a novel combination of a gate function and threshold-triggered synaptic model that are introduced to handle non-differentiability of spikes. In this work, depolarization of membrane potential within a narrow active zone below the firing threshold also induces graded postsynaptic current.

[3] proposes a new type of SNNs, Long Short-Term Memory Spiking Neural Networks (LSNNs) with adapting neurons and support for learning to learn, trained with BPTT with surrogate gradient, demonstrating very good results.

[4] factorizes the standard BPTT into a new form, and proposes three very interesting ideas of converting BPTT into more biologically plausible online learning: (1) an online method to approximate feedback errors, (2) a separate error prediction module trained in the outer loop over a family of different tasks, (3) synthetic gradients combined with eligibility traces for more accurate approximation of the error gradients.

Tempotron uses a "gradient-descent" dynamics and targets only learning timing-based decisions by *single* neurons.

We have a different focus. Superspike [1] is a BP method with surrogate gradient while we more precisely compute gradients through inter and intra dependencies at spiking times. [2] formulates BP at the level of continuous postsynaptic level without directly involving spike timing, which is our focus. In [2], if the membrane potential falls within delta below the firing threshold (activation zone), a graded post-synaptic current will be generated. Differently, we directly consider the all-or-none characteristics of firing spikes. [3] proposes a new recurrent SNN/learning-to-learn network architecture and [4] focuses on the higher-level problem of biologically-plausible online learning. In contrast, we deal with the fundamental problem of BP training with more precisely computed error gradients.

Implementation on neurmorophic hardware: Our TSSL-BP is not biologically plausible and may complicate the implementation on neuromorphic hardware - a limitation. It can train SNNs with high accuracy and low-latency. Low-latency would mitigate its complexity on neuromorphic hardware to a certain extent.

Dynamics over a short time window: We use a short time window of 5 steps to demonstrate the precision of TSSL-BP under low-latency. For most input examples, each trained SNN produces the targeted temporally-varying firing sequences at the output layer. These SNNs are not Time-To-First-Spike networks; neurons are allowed to fire multiple times. Most of the neurons either fire after the first time point or have multiple spikes. Unlike binary ANNs, the trained SNNs here are dynamical. In one SNN, about 20% of neurons fire more than once, 9% of neurons fire more than twice, and 4% of neurons fire more than thrice. We'll include more specific firing statistics in our revised paper.

Intra/inter-neuron dependencies: As in 3.3.2, we split the derivative of a PSC w.r.t a presynaptic spike time $\frac{\partial a[t_k]}{\partial t_m}$ into two parts. First, the spike at t_m directly affects $a[t_k]$, which is called inter dependency. Second, the spike at t_m also affects the succeeding presynaptic spike t_p through resetting which further affects $a[t_k]$. This secondary effect is called intra dependency. Inter-neuron dependencies are dominant in the overall gradients; including the intra-neuron part further improves performance/training speed. Including intra-dependencies in TSSL-BP boosts accuracy by 1.5% for DVS Gesture dataset (40 epochs) and by 4% for CIFAR10 DVS dataset (trained for 5 epochs due to time limitation).

Kernel in loss function: TSSL-BP is flexible about how the loss is defined. The difference between the actual output/targeted firing sequences can be defined via direct comparison, e.g. (6) in the main text, or by using a kernel to measure the so-called Van Rossum distance. The two losses lead to a small performance difference of $< 0.1\%$ for MNIST. Using a kernel to define the loss only smooths the loss but not the firing spikes in the SNN so that the problem of non-differentiable spikes still exists in BPTT with surrogate gradient. Synaptic kernel describes synaptic dynamics and is for a different purpose than the kernel used in the loss. We happen to make the two kernels identical.

Time derivative of membrane potential: As in (3), $\frac{\partial u_i[t_m]}{\partial t_m}$ measures the slope of the membrane potential around firing time t_m. $\frac{\partial u_i[t_m]}{\partial t_m}$ is computed right before the firing: $\frac{\partial u_i[t_m]}{\partial t_m} = \lim_{\Delta t \to 0} \frac{u_i[t_m] - u_i[t_m - \Delta t]}{\Delta t}$ w/o involving thresholding.

[1] Zenke, Friedemann, and Surya Ganguli. "Superspike: Supervised learning in multilayer spiking neural networks."

[2] Huh, Dongsung, and Terrence J. Sejnowski. "Gradient descent for spiking neural networks."

[3] Bellec, Guillaume, et al. "Long short-term memory and learning-to-learn in networks of spiking neurons."

[4] Bellec, Guillaume, et al. "Biologically inspired alternatives to backpropagation through time for learning in recurrent neural nets."