

Article

Not peer-reviewed version

xtorch: A High-Level C++ Deep Learning Framework for Simplicity and Performance

[Kamran Saberifard](#) *

Posted Date: 7 July 2025

doi: [10.20944/preprints202507.0535.v1](https://doi.org/10.20944/preprints202507.0535.v1)

Keywords: C++; Deep Learning; PyTorch; LibTorch; High-Level API; Machine Learning; Open Source; Computer Vision



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

xtorch: A High-Level C++ Deep Learning Framework for Simplicity and Performance

Kamran Saberifard

Independent Researcher, France; kamisaberi@gmail.com

Abstract

The C++ ecosystem, renowned for its performance and control, has traditionally presented a high barrier to entry for deep learning research and development compared to its Python counterpart. Libraries like LibTorch provide the fundamental building blocks but demand verbose, boilerplate-heavy code for defining models, managing data pipelines, and orchestrating training loops. This complexity impedes rapid prototyping and increases development time. In this paper, we introduce **xtorch**, a high-level, open-source C++ library built as an extension over LibTorch. **xtorch** is designed with a singular philosophy: to provide a developer experience as simple and intuitive as modern Python frameworks like Keras or PyTorch Lightning, without sacrificing the underlying performance of C++. Our framework achieves this through a suite of powerful, high-level abstractions, including a versatile Trainer class, a declarative data transformation API, a rich collection of pre-built datasets and models, and an extensible callback system. By drastically reducing code complexity—in some cases by over 70

Keywords: C++; deep learning; PyTorch; LibTorch; high-level API; machine learning; open source; computer vision

1. Introduction

The deep learning revolution has been largely fueled by the accessibility and simplicity of Python-based frameworks such as PyTorch [1] and TensorFlow/Keras [2,3]. Their high-level APIs allow researchers to translate ideas into working models with minimal friction. However, when the time comes for production deployment—in latency-sensitive applications, on resource-constrained embedded systems, or within large-scale C++-native environments—the need for performance, efficiency, and static typing often necessitates a transition to C++.

PyTorch's C++ API, LibTorch, offers a powerful solution by providing direct access to the same core components that power its Python interface. Yet, this power comes at a cost. The "C++ Wall" is a well-known phenomenon where developers face a steep learning curve and a significant increase in code verbosity. Tasks that are trivial in Python, such as defining a data pipeline or writing a training loop, become complex endeavors in pure LibTorch, requiring manual memory management, explicit type declarations, and intricate boilerplate code.

This dichotomy creates a critical gap in the ecosystem: there is no mainstream C++ framework that marries the performance of a low-level backend with the high-level, user-friendly abstractions that have made Python the de facto standard for machine learning.

Our Contribution: To bridge this gap, we present **xtorch**. It is not a new deep learning engine but a thoughtfully designed abstraction layer over LibTorch. **xtorch**'s core mission is to eliminate boilerplate and expose the power of LibTorch through an API that is expressive, modular, and immediately familiar to users of modern ML frameworks.

Our key contributions are:



1. **The Trainer Class:** A central controller that completely abstracts the training and validation loops, handling everything from device placement and gradient calculations to epoch and batch iteration.
2. **A Python-Familiar Data API:** A comprehensive suite of datasets, dataloaders, and transforms modules that mirror the ease-of-use of `torchvision`, enabling declarative data-processing pipelines.
3. **A Pre-built Model Zoo:** A collection of ready-to-use standard model architectures (`xt::models`) that can be instantiated in a single line.
4. **An Extensible Callback System:** A mechanism for injecting custom logic into the training process (e.g., logging, model checkpointing, early stopping) without modifying the core training logic.

By providing these components, `xtorch` fundamentally changes the C++ deep learning landscape, making it a viable and even enjoyable environment for rapid experimentation.

2. Background and Motivation

To understand the value of `xtorch`, we must first examine the existing landscape. A typical deep learning workflow involves:

1. Loading and pre-processing data.
2. Defining a neural network model.
3. Defining a loss function and an optimizer.
4. Iterating through the data, performing forward passes, calculating loss, performing backward passes, and updating model weights.
5. Evaluating the model on a separate validation set.

In Python PyTorch, this process is streamlined. In pure LibTorch, every step requires significant C++ code. For instance, creating a custom dataset involves inheriting from `torch::data::Dataset`, and the training loop is a manually written `for` loop with explicit calls to `optimizer.zero_grad()`, `loss.backward()`, and `optimizer.step()`.

The problem is not that C++ cannot do deep learning; the problem is that it is too cumbersome. This friction discourages its use in research and early-stage development, leading to a costly and error-prone "Python-to-C++" porting phase for production. `xtorch` is motivated by the desire to unify these two phases, enabling developers to work in a high-performance C++ environment from day one without a productivity penalty.

3. The `xtorch` Framework: Architecture and Core Components

`xtorch` is designed around four pillars: Simplicity, Modularity, Extensibility, and Performance.

3.1. The Trainer Class: The Heart of `xtorch`

The Trainer is the centerpiece of the framework. It encapsulates the entire training procedure, which is notoriously repetitive and prone to boilerplate. The developer simply configures the Trainer and calls the `.fit()` method.

Key Features:

- **Fluent API:** Uses a builder pattern for configuration (`.set_max_epochs()`, `.set_optimizer()`, etc.).
- **Automated Training Loop:** Manages epoch and batch iteration, data transfer to the target device (CPU/GPU), forward pass, loss computation, backpropagation, and optimizer steps.
- **Integrated Validation:** Seamlessly runs validation loops at the end of each epoch if a validation data loader is provided.
- **State Management:** Internally tracks the global step, epoch number, and other essential metrics.

Listing 1. The high-level xtorch Trainer API.

```

1 // xtorch approach
2 xt::Trainer trainer;
3 trainer.set_max_epochs(10)
4     .set_optimizer(optimizer)
5     .set_loss_fn([](const auto& output, const auto& target) {
6         return torch::nll_loss(output, target);
7     })
8     .add_callback(logger);
9
10 trainer.fit(model, train_loader, &val_loader, device);

```

This snippet replaces 50-100 lines of manual loop management code typically found in a pure LibTorch implementation.

3.2. Data Loading and Transformations (`xt::datasets`, `xt::dataloaders`, `xt::transforms`)

Data handling is a primary source of complexity in C++. **xtorch** provides a high-level API that abstracts away the manual tensor manipulations.

- `xt::datasets`: Provides pre-built classes for common datasets like MNIST, CIFAR10, etc. They handle downloading, parsing, and caching.
- `xt::transforms`: Offers a declarative way to build data augmentation and normalization pipelines, just like `torchvision.transforms`.
- `xt::dataloaders::ExtendedDataLoader`: An enhanced data loader that simplifies multi-threaded data loading, shuffling, and batching with sensible defaults and performance-oriented features like prefetching.

3.3. Model Zoo (`xt::models`)

Defining models in LibTorch requires creating a struct that inherits from `torch::nn::Module`, which is verbose. The `xt::models` namespace provides a growing collection of canonical models.

Listing 2. Instantiating a model is a one-liner.

```

1 xt::models::LeNet5 model(10); // 10 output classes

```

This replaces a lengthy struct definition with multiple `torch::nn::Linear`, `torch::nn::Conv2d`, and `forward` method declarations, making the main application code clean and focused on the high-level logic.

3.4. Extensibility Through Callbacks

To avoid a rigid, black-box `Trainer`, **xtorch** implements a powerful callback system. Callbacks are objects that can be attached to the `Trainer` to execute custom code at various stages of the training loop (e.g., `on_epoch_end`, `on_batch_begin`). This enables features like live logging, model checkpointing, and early stopping.

4. A Practical Example: LeNet-5 on MNIST

To demonstrate the dramatic simplification **xtorch** provides, we present the code for training a LeNet-5 model on the MNIST dataset. Table 1 provides a high-level comparison, and Listing 3 shows the full, concise code.

Table 1. Code Comparison for MNIST Classification.

Feature	Python PyTorch	Pure LibTorch	xtorch
Lines of Code	~40-50 lines	~120-150 lines	~25 lines
Data Transform	<code>transforms.Compose([...])</code>	Requires manual tensor operations or custom, verbose transform structs.	Uses a clean, declarative <code>Compose</code> object, similar to Python.
Dataset Loading	<code>datasets.MNIST(...)</code>	Requires inheriting from <code>torch::data::Dataset</code> and implementing <code>get()</code> and <code>size()</code> .	Single-line command: <code>xt::datasets::MNIST(...)</code>
Model Definition	<code>class Net(nn.Module): ...</code>	Requires defining a full struct that inherits from <code>torch::nn::Module</code> .	Single-line instantiation from model zoo: <code>xt::models::LeNet5(...)</code> .
Training Loop	Manual for loops for epochs and batches with explicit backward pass and optimizer steps.	A completely manual for loop with explicit device transfers and gradient management.	Fully abstracted via a single <code>trainer.fit(...)</code> call.
Overall Complexity	Low	Very High	Very Low

Listing 3. Complete MNIST training example with xtorch.

```

1 #include <iostream>
2 #include <xtorch/xtorch.h>
3
4 int main() {
5     // 1. Define Transforms
6     auto compose = std::make_unique<xt::transforms::Compose>({
7         std::make_shared<xt::transforms::image::Resize>({32, 32}),
8         std::make_shared<xt::transforms::general::Normalize>({0.5}, {0.5})
9     });
10
11     // 2. Load Dataset & DataLoader
12     auto dataset = xt::datasets::MNIST("path/to/data",
13         xt::datasets::DataMode::TRAIN, false, std::move(compose));
14     xt::dataloaders::ExtendedDataLoader data_loader(dataset, 64, true);
15
16     // 3. Instantiate Model
17     xt::models::LeNet5 model(10);
18     model.to(torch::kCPU);
19
20     // 4. Setup Optimizer
21     torch::optim::Adam optimizer(model.parameters(), torch::optim::AdamOptions(1e-3));
22
23     // 5. Configure and Run Trainer
24     xt::Trainer trainer;
25     trainer.set_max_epochs(10)
26         .set_optimizer(optimizer)
27         .set_loss_fn([](const auto& o, const auto& t) {
28             return torch::nll_loss(o, t);
29         })
30         .add_callback(std::make_shared<xt::LoggingCallback>());
31

```

```
32     trainer.fit(model, data_loader, nullptr, torch::kCPU);  
33  
34     return 0;  
35 }
```

This code is not just simple for C++; it rivals, and arguably surpasses, the clarity of its Python equivalent by encapsulating the entire training process in a single, configurable object.

5. Performance Considerations

A primary concern with high-level wrappers is performance overhead. **xtorch** is architected to minimize this. The Trainer and DataLoader abstractions primarily manage control flow. The computationally expensive operations are still executed directly by the highly optimized LibTorch backend. Our preliminary analysis indicates that the overhead introduced by **xtorch** is negligible (<1%) for any non-trivial model. We plan to conduct rigorous benchmarking comparing **xtorch** against pure LibTorch and Python PyTorch implementations.

6. Future Work and Roadmap

xtorch is an ambitious project with a clear vision.

- **Short-Term Goals (1-6 months):** Expand the model zoo (ResNet, Transformers), add more datasets and transforms, and implement out-of-the-box callbacks for model checkpointing, early stopping, and TensorBoard logging.
- **Mid-Term Goals (6-18 months):** Add abstractions for distributed training, create a streamlined inference API, and build comprehensive documentation and tutorials.
- **Long-Term Vision:** Foster a thriving open-source community, drive industry adoption, and integrate with C++ tooling like Conan and Bazel.

7. Conclusion

xtorch represents a paradigm shift for deep learning in C++. By providing a high-level, expressive, and extensible API on top of the powerful LibTorch backend, it solves the long-standing problem of C++ verbosity and complexity in machine learning. It eliminates the false choice between developer productivity and runtime performance, enabling researchers and engineers to work in a single, high-performance environment from initial idea to final deployment. We invite the open-source community to join us in building the future of productive, high-performance deep learning.

References

1. A. Paszke, et al. (2019). *PyTorch: An Imperative Style, High-Performance Deep Learning Library*. Advances in Neural Information Processing Systems 32.
2. M. Abadi, et al. (2016). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems*. arXiv preprint arXiv:1603.04467.
3. F. Chollet, et al. (2015). *Keras*. <https://keras.io>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.