

ECE 2400 Computer Systems Programming, Fall 2021

PA5: Handwriting Recognition Systems – Tree vs. Table Search

School of Electrical and Computer Engineering
Cornell University

revision: 2021-11-28-16-44

This handout is written assuming students must implement the tree search, table search, and an alternative handwriting recognition system. For 2021, we only require students to implement *either* tree search *or* table search. Students can choose which of these two systems to implement. If students choose not to implement tree search then they do not need to implement `Tree<T, CmpFunc>`. If students choose not to implement table search then they do not need to implement `Table<T, HashFunc>`. All students must implement an alternative handwriting recognition system.

1. Introduction

The fifth programming assignment will enable you to build off of the work you did in the previous programming assignment to continue to apply all of the concepts you have learned throughout the semester in the context of a machine learning system. You will continue to leverage what you have learned on algorithms (e.g., iteration vs. recursion), data structures (e.g., vectors), complexity analysis, C++ basics (e.g., namespaces, references, exceptions, and dynamic allocation), and object-oriented programming (e.g., classes, member functions, constructors, operator overloading, the rule of three, data encapsulation, interface vs. implementation). You will also now leverage what you have learned on generic programming (e.g., templates), functional programming, tree data structures, and table data structures. You will apply this understanding to implement, test, and evaluate a handwriting recognition system that can classify handwritten numbers into ten classes, the digits from zero through nine, with high accuracy.

In this assignment, you will revisit the supervised learning models that classify handwritten digits developed in the previous programming assignment, and you will add three new models. Recall that these models have two main phases: *training* and *classification*. In the *training* phase, the model is provided with a large set of images, each with a label (e.g., '1', '7', '9') that indicates the corresponding digit. In the *classification* phase, the model is provided with new inputs that it has never seen before and predicts their label based on what it has learned in the training phase. You will use the classic MNIST database of handwritten digits. The database is composed of 70,000 examples of images each with 28×28 grayscale pixels. Each pixel has a value between 0 and 255 where 0 represents white, 255 represents black, and intermediate values represent intermediate levels of gray. The dataset is divided into a training dataset of 60,000 images and a test dataset of 10,000 images for evaluation. Figure 1 shows a few images from the MNIST dataset. The digits were handwritten by several hundred different writers ranging from average high school students to Census Bureau employees. This means that the legibility of the handwritten digits varies significantly. The goal of the assignment is to design a system that can classify images of these handwritten digits with high accuracy into one of ten classes: the numbers zero through nine. You will start off by first implementing a generic sorting algorithm before implementing four classes: `Vector<T>`, `Image`, `Tree<T, CmpFunc>`, and `Table<T, HashFunc>`. Then you will leverage these classes to reimplement the two simple classification algorithms from the previous programming assignment (one based on linear search and the other based on binary search), before implementing two new classification algorithms (one based on



Figure 1: Four Example MNIST Images – Images include 28×28 grayscale pixels and a label. Each pixel has a value between 0 and 255 where 0 represents white, 255 represents black, and intermediate values represent intermediate levels of gray.

tree search and the other based on table search). Finally, you will have the freedom to implement an alternative supervised learning model using the standard C++ library, external libraries, and whatever algorithm you choose. As in the previous assignments, we will leverage the CMake framework for building programs, the CTest framework for unit testing, GitHub Actions for continuous integration testing, and lcov for code coverage analysis. You will also have the opportunity to combine a Python GUI front-end with your C++ handwriting recognition back-end to create a complete system which will enable users to draw numbers which can then be (hopefully correctly) classified!

After your handwriting recognition systems are functional and tested, you will evaluate the accuracy and performance trade-offs between implementations. You should include all of the optimizations you explored in the previous programming assignment. You will write a report that includes your complexity analysis, a discussion of your optimizations, and a quantitative evaluation of the performance across all implementations. **You should consult the programming assignment logistics document for more information about the expectations for all programming assignments and how they will be assessed. While the final code and report are all due at the end of the assignment, we also require meeting an incremental milestone in this PA. Requirements specific to this PA for the incremental milestone and the final report are described at the end of this handout.**

This handout assumes that you have read and understand the course tutorials and that you have attended the discussion sections. To get started, log in to an ecelinux server, source the setup script, and clone your individual remote repository from GitHub:

```
% source setup-ece2400.sh
% mkdir -p ${HOME}/ece2400
% cd ${HOME}/ece2400
% git clone git@github.com:cornell-ece2400/groupid
% cd ${HOME}/ece2400/groupid/pa5-sys
% tree
```

Where `groupid` should be replaced with your group number. **You should never fork your individual remote repository! If you need to work in isolation then use a branch within your individual remote repository.** If you have already cloned your individual remote repository, then use `git pull` to ensure you have any recent updates before working on your programming assignment.

```
% cd ${HOME}/ece2400/groupid
% git pull
% tree pa5-sys
```

For this assignment, you will work in the `pa5-sys` subproject, which includes the following files:

- `CMakeLists.txt` – CMake configuration script to gen Makefile
- `src/ece2400-stdlib.h` – Header file for course standard library

- `src/ece2400-stdlib.cc`
- `src/ece2400-stdlib.inl`
- `src/sort.h`
- `src/sort.inl`
- `src/sort-adhoc.cc`
- `src/Vector.h`
- `src/Vector.inl`
- `src/vector-adhoc.cc`
- `src/Image.h`
- `src/Image.cc`
- `src/Image.inl`
- `src/image-adhoc.cc`
- `src/Tree.h`
- `src/Tree.inl`
- `src/tree-adhoc.cc`
- `src/Table.h`
- `src/Table.inl`
- `src/table-adhoc.cc`
- `src/IHandwritingRecSys.h`
- `src/HRSLinearSearch.h`
- `src/HRSLinearSearch.cc`
- `src/HRSBinarySearch.h`
- `src/HRSBinarySearch.cc`
- `src/HRSTreeSearch.h`
- `src/HRSTreeSearch.cc`
- `src/HRSTableSearch.h`
- `src/HRSTableSearch.cc`
- `src/HRSAAlternative.h`
- `src/HRSAAlternative.cc`
- `src/digits.dat`
- `test/image-directed-test.cc`
- `test/image-random-test.cc`
- `test/sort-directed-test.h`
- `test/sort-random-test.h`
- `test/sort-int-directed-test.cc`
- `test/sort-int-random-test.cc`
- `test/sort-image-directed-test.cc`
- `test/sort-image-random-test.cc`
- `test/vector-directed-test.h`
- `test/vector-random-test.h`
- `test/vector-int-directed-test.cc`
- `test/vector-int-random-test.cc`
- `test/vector-image-directed-test.cc`
- `test/vector-image-random-test.cc`
- `test/tree-directed-test.h`
- `test/tree-random-test.h`
- `test/tree-int-directed-test.cc`
- Source code for course standard library
- Inline code for course standard library
- Header file for `sort<T, CmpFunc>`
- Inline source code for `sort<T, CmpFunc>`
- Ad-hoc test program for `sort<T, CmpFunc>`
- Header file for `Vector<T>`
- Inline source code for `Vector<T>`
- Ad-hoc test program for `Vector<T>`
- Header file for `Image`
- Source code for `Image`
- Inline source code for `Image`
- Ad-hoc test program for `Image`
- Header file for `Tree<T, CmpFunc>`
- Inline source code for `Tree<T, CmpFunc>`
- Ad-hoc test program for `Tree<T, CmpFunc>`
- Header file for `Table<T, HashFunc>`
- Inline source code for `Table<T, HashFunc>`
- Ad-hoc test program for `Table<T, HashFunc>`
- Header file for the HRS interface class
- Header file for `HRSLinearSearch`
- Source code for `HRSLinearSearch`
- Header file for the HRS `HRSBinarySearch`
- Source code for the HRS `HRSBinarySearch`
- Header file for the HRS `HRSTreeSearch`
- Source code for the HRS `HRSTreeSearch`
- Header file for the HRS `HRSTableSearch`
- Source code for the HRS `HRSTableSearch`
- Header file for the HRS `HRSAAlternative`
- Source code for the HRS `HRSAAlternative`
- Data file with selected test images
- Directed test cases for `Image`
- Random test cases for `Image`
- Directed test cases for `sort<T, CmpFunc>`
- Random test cases for `sort<T, CmpFunc>`
- Directed test cases for `sort<int>`
- Random test cases for `sort<int>`
- Directed test cases for `sort<Image>`
- Random test cases for `sort<Image>`
- Directed test cases for `Vector<T>`
- Random test cases for `Vector<T>`
- Directed test cases for `Vector<int>`
- Random test cases for `Vector<int>`
- Directed test cases for `Vector<Image>`
- Random test cases for `Vector<Image>`
- Directed test cases for `Tree<T, CmpFunc>`
- Random test cases for `Tree<T, CmpFunc>`
- Directed test cases for `Tree<int, IntCmpFunc>`

- test/tree-int-random-test.cc – Random test cases for Tree<int, IntCmpFunc>
- test/tree-image-directed-test.cc – Directed test cases for Tree<Image, ImgCmpFunc>
- test/tree-image-random-test.cc – Random test cases for Tree<Image, ImgCmpFunc>
- test/table-directed-test.h – Directed test cases for Table<T, HashFunc>
- test/table-random-test.h – Random test cases for Table<T, HashFunc>
- test/table-int-directed-test.cc – Directed test cases for Table<int, IntHashFunc>
- test/table-int-random-test.cc – Random test cases for Table<int, IntHashFunc>
- test/table-image-directed-test.cc – Directed test cases for Table<Image, ImgHashFunc>
- test/table-image-random-test.cc – Random test cases for Table<Image, ImgHashFunc>

- test/hrs-linear-search-directed-test.cc – Directed test cases for HRSLinearSearch
- test/hrs-binary-search-directed-test.cc – Directed test cases for HRSBinarySearch
- test/hrs-tree-search-directed-test.cc – Directed test cases for HRSTreeSearch
- test/hrs-table-search-directed-test.cc – Directed test cases for HRSTableSearch
- test/hrs-alternative-directed-test.cc – Directed test cases for HRSAAlternative

- eval/hrs-linear-search-eval.cc – Evaluation Program for HRSLinearSearch
- eval/hrs-binary-search-eval.cc – Evaluation Program for HRSBinarySearch
- eval/hrs-tree-search-eval.cc – Evaluation Program for HRSTreeSearch
- eval/hrs-table-search-eval.cc – Evaluation Program for HRSTableSearch
- eval/hrs-alternative-eval.cc – Evaluation Program for HRSAAlternative
- eval/hrs-backed.cc – Backend for interactive classification system
- eval/hrs-frontend.py – Frontend for interactive classification system

By the end of this programming assignment you will have written over a thousand lines of code, and you will be able to read and understand almost ten thousand lines of code which are included in this assignment spanning implementation, testing, and evaluation.

The programming assignment is divided into the following steps. Complete each step before moving on to the next step.

- Step 1. Implement and test generic sort<T, CmpFunc> function
- Step 2. Implement and test generic Vector<T> class
- Step 3. Implement and test Image class
- Step 4. Implement and test generic Tree<T, CmpFunc> class
- Step 5. Implement and test generic Table<T, HashFunc> class
- Step 6. Implement and test HRSLinearSearch class
- Step 7. Implement and test HRSBinarySearch class
- Step 8. Implement and test HRSTreeSearch class
- Step 9. Implement and test HRSTableSearch class
- Step 10. Implement and test HRSAAlternative class

We cannot stress enough how important it is to take an incremental design approach! You really must implement *and test* each step before trying to implement the next step. This means more than just adhoc testing. You must do thorough directed and random testing of each step *before* implementing the next step.

2. Implementation Specifications

The high-level goal for this programming assignment is to implement a handwriting recognition system with five classification algorithms. You will start off by implementing a generic sorting algorithm (sort<T, CmpFunc>) which is basically a generic version of the sorting algorithm you im-

plemented in the previous programming assignment. The sorting algorithm should work for any type `T` and also takes a comparison function object. You will then implement a generic `Vector<T>` data structure which is basically a generic version of the resizable vector you implemented in the previous programming assignment. Your `Vector<T>` will make use of `sort<T,CmpFunc>`. `Image` uses a `Vector<int>` to store pixels, and you can also now use `Vector<Image>` to create a vector that stores `Images`. You will also implement a generic tree data structure (`Tree<T,CmpFunc>`) and table data structure (`Table<T,HashFunc>`). You will then leverage these data structures to construct the four handwriting recognition systems: `HRSLinearSearch`, `HRSBinarySearch`, `HRSTreeSearch`, and `HRSTableSearch`. You will then implement a `HRSAAlternative` system which can use any approach you like.

Note that your implementations cannot use anything from the Standard C library except for the `printf` function defined in `stdio.h`, the `MIN/MAX` macros defined in `limits.h`, the `NULL` macro defined in `stddef.h`, and the `assert` macro defined in `assert.h`. Your implementations cannot use anything from the Standard C++ library except for C++ I/O streams from `iostream` and `std::function` from `functional`. **As an exception, your alternative handwriting recognition system can use anything in the standard C++ library and/or external libraries.**

2.1. `sort<T,CmpFunc>` Algorithm

You will implement a generic sort algorithm that can sort an array of elements of any type `T`. This algorithm is similar to what you implemented in the previous programming assignment; the difference is that it should be generic across any type `T`, and that it takes a comparison function object (i.e., function pointer, functor, lambda) which should be used by your algorithm to compare different objects of type `T`. The interface for the generic sorting function is as follows:

```
template < typename T, typename CmpFunc >
void sort( T* a, int size, CmpFunc cmp );
```

Note how we use the `CmpFunc` template parameter to indicate the type of the comparison function `cmp`. More specifically, the comparison function `cmp` should be callable with the following function signature: `bool cmp(const T& a, const T& b)`. The function should return `true` if object `a` is strictly less than `b`, and `false` otherwise. In this PA, there is no guarantee that comparison operators (e.g., `operator<`) have been defined for type `T`, so your sorting algorithm should use `cmp(a,b)` in places where you want to compare the value of `a` and `b`. For example, `a < b` is simply `cmp(a,b)`, `a > b` is equivalent to `cmp(b,a)`, and `a == b` is equivalent to `!cmp(a,b) && !cmp(b,a)`. Your algorithm must work correctly if `size` is zero, which means the input array pointer may be a `nullptr`.

The interface for `sort<T,CmpFunc>` is provided for you in `src/sort.h`. Write the implementation of your function inside of `src/sort.inl`. Note that since this is a templated data structure, all of the templated definitions must be placed in the `.inl` file not in a `.cc` file. We cannot compile the function template since we don't know the types of `T` nor `CmpFunc` yet! We can only compile a function template specialization.

2.2. `Vector<T>` Data Structure

You will implement a generic resizable vector data structure that stores data of type `T`. This data structure is similar to what you implemented in the previous programming assignment; the differences are: (1) it should be generic across any type `T`; (2) its `find_closest` member function is templated and takes a distance function object as an argument; and (3) its `sort` member function is templated

and takes a comparison function object as an argument. You are responsible for implementing each of the following functions except for `print` which is already implemented:

```

Vector<T>::Vector();
Vector<T>::Vector( T* array, int size );

Vector<T>::~~Vector();
Vector<T>::Vector( const Vector<T>& vec );
Vector<T>& Vector<T>::operator=( const Vector<T>& vec );

int      Vector<T>::size() const;
void     Vector<T>::push_back( const T& value );
const T& Vector<T>::at( int idx ) const;
T&      Vector<T>::at( int idx );
bool     Vector<T>::contains( const T& value ) const;

template <typename DistFunc>
T        Vector<T>::find_closest_linear( const T& value, DistFunc dist ) const;

template <typename DistFunc, typename CmpFunc>
T        Vector<T>::find_closest_binary( const T& value, int k,
                                         DistFunc dist, CmpFunc cmp ) const;

template <typename CmpFunc>
void     Vector<T>::sort( CmpFunc cmp );

void     Vector<T>::print() const;
const T& Vector<T>::operator[]( int idx ) const;
T&      Vector<T>::operator[]( int idx );

```

The specification for these functions is as follows:

- `Vector<T>::Vector()`
The default constructor for `Vector<T>`. It constructs an empty `Vector<T>` by initializing all member fields in `Vector<T>`.
- `Vector<T>::Vector(int* arr, int size)`
A non-default constructor that constructs a `Vector<T>` from an array of `T`s with the given `size`. This constructor should perform a deep copy, i.e., copy the values inside the array rather than copying the pointer. Modifying or deleting the input array after construction should have no effect on the constructed `Vector<T>` object. You can assume that the input `size` is never greater than the actual size of the array `arr`. Construct an empty `Vector<T>` if `size` is 0.
- `Vector<T>::~~Vector()`
Destructor for `Vector<T>`. It frees the memory allocated on the heap by `Vector<T>`. Note that you should use the `delete` and `delete[]` operator instead of `free` as in C.
- `Vector<T>::Vector(const Vector<T>& vec)`
Copy constructor that performs a deep copy. It should copy the values stored in `vec`. Subsequent actions on `vec` should have no effect on the constructed `Vector<T>`. *You must carefully handle the case when copying from an empty vector!*

- `Vector<T>& Vector<T>::operator=(const Vector<T>& vec)`
This overloads the assignment operator. Copy the values stored in `vec`. Note that if the current `Vector<T>` is not empty, you need to free the memory allocated for it first. *It is very important that you carefully handle the case of self assignment! You must also carefully handle the case when assigning from an empty vector!*
- `int Vector<T>::size() const`
Return the current number of elements in the vector. If the `Vector<T>` is empty, this function should return 0.
- `void Vector<T>::push_back(const T& value)`
Push a new element with the given value `value` onto the end of the `Vector<T>`. If there is not enough allocated space, dynamically allocate more memory to store both existing elements and the new element. Note that you should use the `new` or `new[]` operator instead of `malloc` as in C.
- `const T& Vector<T>::at(int idx) const`
Return the value at the given index `idx` of the `Vector<T>`. If the given index is out-of-bound, throw `ece2400::OutOfRange` with a useful error message. Note that this is the “read” version of `at` since it is declared `const` and returns a `const` reference.
- `T& Vector<T>::at(int idx)`
Return the value at the given index `idx` of the `Vector<T>`. If the given index is out-of-bound, throw `ece2400::OutOfRange` with a useful error message. Note that this is the “write” version of `at` since it returns a non-`const` reference. It enables writing items in the `Vector<T>`.
- `bool Vector<T>::contains(const T& value) const`
Search the `Vector<T>` for the given value and returns `true` if the value is found and `false` otherwise. If the `Vector<T>` is empty, then this function should just return `false`.
- `template <typename DistFunc>`
`T Vector<T>::find_closest_linear(const T& value, DistFunc dist) const`
Perform a linear search of the vector for the value that is closest to the given value (`value`), as measured by the `dist` distance function object. The distance function object `dist` should be callable with the following function signature: `int dist(const T& a, const T& b)`. The function should take two inputs of type `T` and return their distance as `int`. If multiple values have the same minimal distance, return the one that has the lowest index. As in the previous PA, is undefined if calculating the distance results in overflow. If the `Vector<T>` is empty, throw `ece2400::OutOfRange` with proper error message. Note that since this is a template member function in a template class, you need the following syntax to define `Vector<T>::find_closest_linear`:

```
template <typename T>
template <typename DistFunc>
T Vector<T>::find_closest_linear( const T& value, DistFunc dist ) const
{
    // ...
}
```
- `template <typename DistFunc, typename CmpFunc>`
`T Vector<T>::find_closest_binary(const T& value, int k,`
`DistFunc dist, CmpFunc cmp) const`
Perform a binary search of the vector to find an index that has a value close to the given value (`value`), and then perform a linear search of `K` elements centered around the index determined by the binary search. The binary search should use the given `cmp` comparison function object,

while the linear search should use the given `dist` comparison function. As in the previous PA, you will search $K/2$ items backwards and $K/2$ items forwards from the item found during the binary search. You will need to carefully handle the case where there are less than $K/2$ items either before or after the item found during the binary search. Return the closest value from the linear search. As in the previous PA, is undefined if calculating the distance results in overflow. If the `Vector<T>` is empty, throw `ece2400::OutOfRange` with proper error message. You must check that the `Vector<T>` is sorted before doing the binary search, and throw `ece2400::InvalidArgument` with a proper error message if it is not sorted.

```
template <typename T>
template <typename DistFunc, typename CmpFunc>
T Vector<T>::find_closest_binary( const T& value, int k,
                                DistFunc dist, CmpFunc cmp ) const
{
    // ...
}
```

- `template <typename CmpFunc>`
`void Vector<T>::sort(CmpFunc cmp)`

Sort the internal array in an ascending order defined by the comparison function `cmp`. You must use the generic `sort<T, CmpFunc>` function developed in the previous step. Note that since this is a template member function in a template class, you need the following syntax in the definition of `Vector<T>::sort`:

```
template <typename T>
template <typename CmpFunc>
void Vector<T>::sort( CmpFunc cmp )
{
    ::sort( ... );
}
```

Notice now we need to explicitly use `::` so the compiler knows we want to call the generic global `sort` function and not the `sort` member function.

- `void Vector<T>::print() const`
Print the content in the `Vector<T>`. This member function is used for debugging purposes. Note for this member function to be generic it needs to use the C++ `iostream` library (e.g., `std::cout`) instead of `printf`. This is because the type `T` may not even have a format specifier for `printf`. We provide you with this function.
- `const T& Vector<T>::operator[](int idx) const`
This overloads the subscript operator. It should just return the value at the given index `idx` of the vector without any boundary check. Note that this is different from `at` as `at` throws an exception when index is out-of-bound. Note that this is the “read” version since it is declared `const` and returns a `const` reference.
- `T& Vector<T>::operator[](int idx)`
This overloads the subscript operator. It should just return the value at the given index `idx` of the vector without any boundary check. Note that this is different from `at` as `at` throws an exception when index is out-of-bound. Note that this is the “write” version since it returns a non-`const` reference. It enables writing items in the `Vector<T>`.

The functions vary in complexity, and some may require just a few lines of code to implement. To give you an idea of how to use this class, here is a simple function that constructs a `Vector<int>`, pushes back three values, gets the middle value, and then destructs the `Vector<int>`:

```
int main( void )
{
    Vector<int> vec;          // Declare a Vector<int> on the stack
    vec.push_back ( 11 );   // Push back 11
    vec.push_back ( 12 );   // Push back 12
    vec.push_back ( 13 );   // Push back 13
    int a = vec.at ( 1 );   // int a now has 12
}
```

The interface for `Vector<T>` is provided for you in `src/Vector.h`. Write the implementation of your member variables inside `src/Vector.h` and the implementation of each function inside of `src/Vector.inl`. Note that since this is a templated data structure, all of the templated definitions must be placed in the `.inl` file not in a `.cc` file. We cannot compile the class template since we don't know the type `T` yet! We can only compile a class template specialization.

2.3. Image Data Structure

After `Vector<T>` is implemented and tested, you will then use it to implement the `Image` class. `Image` uses `Vector<int>` to store an array of integers. Each integer represents a pixel in grayscale and has a value within the range of `[0, 255]`. Lower numbers represent lighter shades (with 0 representing white), while higher numbers represent darker shades (with 255 representing black). Each `Image` object has a label associated with it. `Image` also has an `intensity`, which we define here as the sum of all pixels. You are responsible for implementing each of the following functions except for `print`, `display`, and the `<<` operator, which have been implemented for you:

```
Image::Image();
Image::Image( const Vector<int>& vec, int ncols, int nrows );

int Image::get_ncols() const;
int Image::get_nrows() const;
int Image::at( int x, int y ) const;
void Image::set_label( char l );
char Image::get_label() const;
int Image::get_intensity() const;
int Image::distance( const Image& other ) const;
void Image::print() const;
void Image::display() const;

bool Image::operator==( const Image& rhs ) const;
bool Image::operator!=( const Image& rhs ) const;

const int& Image::operator[]( int idx ) const;

std::ostream& operator<<( std::ostream& os, const Image& image );
```

Here is a brief specification for each member function of the `Image` class:

- `Image::Image()`
Default constructor for `Image`. This function constructs an empty `Image` by initializing all data members and setting the label to `'?'`.
- `Image::Image(const Vector<int>& vec, int ncols, int nrows)`
Non-default constructor that constructs an `Image` from a `Vector<T>` given the number of columns (`ncols`) and number of rows (`nrows`). If the size of the vector does not match the number of columns and number of rows, throw `ece2400::InvalidArgument` with a useful error message.
- `int Image::get_ncols() const`
Return the number of columns of the current `Image`. Return 0 if the current `Image` is empty.
- `int Image::get_nrows() const`
Return the number of rows of the current `Image`. Return 0 if the current `Image` is empty.
- `int Image::at(int x, int y) const`
Return the value of the pixel at x -th column and y -th row. For example, if an `Image` is constructed from `{0,1,2,3}` with `ncols` and `nrows` both equal to 2, then `at(0,0)` returns 0, `at(1,0)` returns 1, `at(0,1)` returns 2, `at(1,1)` returns 3. If x or y is out-of-bounds, throw `ece2400::OutOfRange` with proper error message.
- `void Image::set_label(char label)`
Set the current label of the `Image` to the given character `label`.
- `char Image::get_label() const`
Return the current label of the `Image`. If no `set_label` has been called, return `'?'`.
- `int Image::get_intensity() const`
Return the intensity of the current `Image`. Note that intensity here is simply defined as the sum of all pixels. You can assume no overflow will occur for all arithmetic operations in this function.
- `int Image::distance(const Image& other)`
Return the square of the Euclidean distance between this image and image `other`, which is just the sum of the difference between each pixel squared. For example, if image `a` has four pixels `{1,9,9,5}` and `b` has four pixels `{0,4,2,3}` then `distance(a,b)` should return $(1 - 0)^2 + (9 - 4)^2 + (9 - 2)^2 + (5 - 3)^2 = 55$. If the dimensions of the two images do not match, throw `ece2400::InvalidArgument` with a useful error message. Since an `Image` cannot be larger than 128×128 and each pixel cannot be larger than 255, it should be possible to write this function without worrying about overflow.
- `Image::print() const`
Prints the label and intensity of the `Image`. We provide you with this function.
- `Image::display() const`
Prints `Image` using the `print_pixel` helper function, which prints a block to the terminal that is a shade of grey determined by the pixel value. We provide you with this function.
- `bool Image::operator==(const Image& rhs) const`
Overload the equal to operator so that it compares the value of each pixel. Return `true` only if the each pixel in the right-hand-side image is the same as that in the current image. If the dimension of the two image does not match, simply return `false`. Otherwise return `true`.
- `bool Image::operator!=(const Image& rhs) const`
Overload the equal to operator so that it compares the value of each pixel. Return `false` only

if the each pixel in the right-hand-side image is the same as that in the current image. If the dimension of the two image does not match, simply return `true`. Return `false` if both images are empty.

- `const int& Image::operator[](int idx) const`
This overloads the subscript operator. It should just return the pixel value at the given index `idx` of the image without any boundary check. Note that this is different from `at` as `at` throws an exception when index is out-of-bound. Note that this is the “read” version since it is declared `const` and returns a `const` reference.
- `std::ostream& operator<<(std::ostream& os, const Image& image);`
Overloads the `<<` operator as a free function (not as a member function) to be able to display an `Image` using `std::cout`. Allows you to use the `<<` operator to insert the image’s intensity and label into the given `os` output stream. We provide you with this function.

The functions vary in complexity, and some may require just a few lines of code to implement. The interface for `Image` is provided for you in `src/Image.h`. Write the implementation of your member variables inside `src/Image.h` and the implementation of each function inside of `src/Image.cc`.

2.4. `Tree<T, CmpFunc>` Data Structure

You will implement a generic binary search tree data structure that stores data of type `T` whose ordering is established by `CmpFunc`. Each internal node in the binary search tree stores a value, a pointer to the left subtree, and a pointer to the right subtree. The value stored in any node should be greater or equal to any value stored in the left subtree and smaller than any value stored in the right subtree. Both subtrees should also recursively satisfy this property. You will need to define a nested struct or nested class which represents a node of the tree inside `Tree<T, CmpFunc>`. Your node should have the following fields: a field to store values of type `T`, a pointer to a node for the left subtree, and a pointer to a node for the right subtree. You are responsible for implementing each of the following functions:

```
Tree<T,CmpFunc>::Tree( unsigned int K, CmpFunc cmp );

Tree<T,CmpFunc>::~~Tree();
Tree<T,CmpFunc>::Tree( const Tree<T,CmpFunc>& tree );
Tree<T,CmpFunc>& Tree<T,CmpFunc>::operator=( const Tree<T,CmpFunc>& tree );

int      Tree<T,CmpFunc>::size() const;
void     Tree<T,CmpFunc>::add( const T& value );
bool     Tree<T,CmpFunc>::contains( const T& value ) const;
Vector<T> Tree<T,CmpFunc>::to_vector() const;

template <typename DistFunc>
T Tree<T,CmpFunc>::find_closest( const T& value, DistFunc dist ) const;

void Tree<T,CmpFunc>::print() const;
```

Here is a brief specification for each member function of `Tree<T, CmpFunc>` class:

- `Tree<T, CmpFunc>::Tree(int K, CmpFunc cmp)`
Non-default constructor for `Tree<T, CmpFunc>`. Note that `K` is a parameter for `find_closest`,

which corresponds to the number of elements for exhaustive search. `cmp` is the comparison function that this tree data structure should use to establish ordering between values of type `T`. This function constructs an empty `Tree<T, CmpFunc>` by initializing all data members, stores the value of `K`, and copies the comparator so that you can call it later.

- `Tree<T, CmpFunc>::~~Tree()`
Destructor for `Tree<T, CmpFunc>`. It frees the memory allocated on the heap by `Tree<T, CmpFunc>`. Note that you should use the `delete` and `delete[]` operator instead of `free` as in C. *Hint: You will likely need to use a private recursive helper function to implement this member function.*
- `Tree<T, CmpFunc>::Tree(const Tree<T, CmpFunc>& tree)`
Copy constructor that performs a deep copy. It should copy the values stored in `tree`. Subsequent actions on `tree` should have no side effect on the constructed `Tree<T, CmpFunc>`. *You must carefully handle the case when copying from an empty tree! Hint: You will likely need to use a private recursive helper function to implement this member function.*
- `Tree<T, CmpFunc>& Tree<T, CmpFunc>::operator=(const Tree<T, CmpFunc>& tree)`
This overloads the assignment operator. Copy the values stored in `tree`. Note that if the current `Tree<T, CmpFunc>` is not empty, you need to free the memory allocated for it first. *It is very important that you carefully handle the case of self assignment! You must carefully handle the case when assigning from an empty tree! Hint: You will likely need to use a private recursive helper function to implement this member function.*
- `int Tree<T, CmpFunc>::size() const`
Return the current number of elements in the tree. If the `Tree<T, CmpFunc>` is empty, this function should return 0. *Hint: While you could use a private recursive helper function to implement this member function, you might want to consider using a constant time implementation.*
- `void Tree<T, CmpFunc>::add(const T& value)`
Add a new node with the given value `value` to the `Tree<T, CmpFunc>`. The new node should be dynamically allocated on the heap. The binary search property for `Tree<T, CmpFunc>` should still hold after adding the new node. You will need to first traverse the tree to find the leaf node to which the new node should be added. We recommend you implement a recursive helper function to do so. If the value is already in the tree, this function should simply return and do nothing. Note that you should use the `new` operator instead of `malloc` as in C. *Hint: You will likely need to use a private recursive helper function to implement this member function.*
- `bool Tree<T, CmpFunc>::contains(const T& value) const`
Search the `Tree<T, CmpFunc>` for the given value and returns `true` if the value is found and `false` otherwise. If the `Tree<T, CmpFunc>` is empty, then this function should just return `false`. *Hint: You will likely need to use a private recursive helper function to implement this member function.*
- `Vector<T> Tree<T, CmpFunc>::to_vector() const`
Return a `Vector<T>` that contains all the items stored in the `Tree<T, CmpFunc>`. The `Vector<T>` should represent an in-order traversal of the tree (i.e., first add all values in the left subtree, then add the value of the current node, and finally add all values in the right subtree). *Hint: You will likely need to use a private recursive helper function to implement this member function.*
- `template <typename DistFunc>`
`const T& Tree<T, CmpFunc>::find_closest(const T& value, DistFunc dist) const`
Search the `Tree<T, CmpFunc>` and return the value that has (approximately) the smallest difference (as defined by the `dist` distance function) from the given value. The distance function takes two inputs of type `T`, and returns the distance between them as an `int`, and has a function

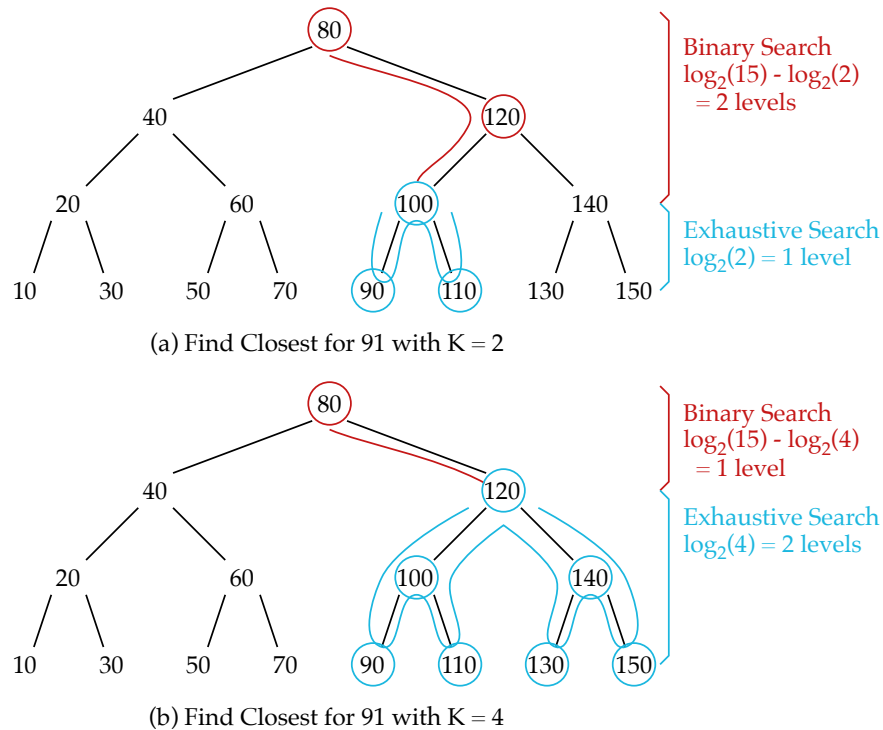


Figure 2: `Tree<T>::find_closest` Example – Example for balanced binary search tree with 15 integers. There are $\log_2(15) = 3$ levels. (a) With $K = 2$, we do a binary search through $\log_2(15) - \log_2(2) = 2$ levels, and then we do an exhaustive search through the remaining sub-tree which in this case has three nodes. (b) With $K = 4$, we do a binary search through $\log_2(15) - \log_2(4) = 1$ level, and then we do an exhaustive search through the remaining sub-tree which in this case has seven nodes. The binary search uses intensity, while the exhaustive search uses Euclidean distance. Larger K improves accuracy while reducing performance.

signature of `int dist(const T& a, const T& b)`. You should use a hybrid of binary search and exhaustive search (see Figure 2). You should perform a partial binary search to get down to a certain level in the tree and then perform an exhaustive search on the subtree. We recommend you first implement the exhaustive search by calling your recursive helper function from `Tree<T, CmpFunc>::to_vector` on the node you want to start your exhaustive search from, and then using `Vector<T>::find_closest_linear` on the vector that is returned. Implementing the exhaustive search directly on the tree is a performance optimization you can explore later. You will need to compute the number of levels for binary search and exhaustive search first. When doing so you can assume the tree is perfectly balanced, thus the total number of levels is $\log_2(N)$ where N is the number of elements in the tree. The number of levels for exhaustive search is $\log_2(K)$, and the number of levels for binary search is simply the total number of levels minus the number of levels for exhaustive search. We have provided you with a `ece2400::log2` function in `ece2400-stdlib.h`. If the `Tree<T, CmpFunc>` is empty, throw `ece2400::OutOfRange` with proper error message. You can assume no overflow will occur for all arithmetic operations in this function. *Hint: You will likely need to use one or more private recursive helper functions to implement this member function.* Note that this function is a template member function in a template class, so you need the following syntax in the function definition:

```
template <typename T, typename CmpFunc>
```

```

template <typename DistFunc>
T Tree<T,CmpFunc>::find_closest( const T& value, DistFunc dist ) const
{
    // ...
}

```

You need a distance function in order to call this member function. You can create a free function, functor, or lambda to calculate the distance. Distance for integers is the difference between the two values (make sure to return a positive number). Distance for images should be the Euclidean distance between the pixel arrays, just like the `Image::distance()` member function in Section 2.3.

- `void Tree<T,CmpFunc>::print() const`
Print the content in the `Tree<T,CmpFunc>`. This member function is used for your own debugging purpose. You can implement this function in any way you like. You do not need to test this function. However, note that to make this member function generic you will need to use the C++ `iostream` library (e.g., `std::cout`) instead of `printf`. We have provided a commented-out version of our implementation of this function, which prints the tree with the root node on the left. You can leverage this implementation by replacing our references to the private member fields of the tree with the names of your private member fields.

2.5. `Table<T,HashFunc>` Data Structure

You will implement a generic table data structure that stores data of type `T` and uses a hash function of type `HashFunc` to map values of type `T` to integers. Internally, it stores data in a vector of vectors (i.e., `Vector<Vector<T>>`). Each `Vector<T>` corresponds to a bin that stores values within a certain range, which is similar to the bucket sort algorithm you implemented in a previous programming assignment. The hash function (type `HashFunc`) associated with the table data structure helps translate objects of unknown type `T` into integers which we can then use as indices into the internal storage. You are responsible for implementing each of the following functions:

```

Table<T,HashFunc>::Table( int k, HashFunc hash );

int      Table<T,HashFunc>::size() const;
void     Table<T,HashFunc>::add( const T& value );
bool     Table<T,HashFunc>::contains( const T& value ) const;
Vector<T> Table<T,HashFunc>::to_vector() const;

template <typename DistFunc>
T Table<T,HashFunc>::find_closest( const T& value, DistFunc dist ) const;

void Table<T,HashFunc>::print() const;

```

Note that you do not need to implement the copy constructor, assignment operator, or the destructor because you will use your `Vector<T>` which already has those implemented. Here is a brief specification for each member function of `Table<T,HashFunc>` class:

- `Table<T,HashFunc>::Table<T,HashFunc>(int k, HashFunc hash)`
Non-default constructor for `Table<T,HashFunc>`. This function constructs an empty `Table<T,HashFunc>` by initializing all data members and copies the hash function. The hash function `hash` should take a value of type `T` and hash it to a positive integer of type `int`. `k` corresponds to the load factor

of the `Table<T,HashFunc>`. The load factor determines when you should increase the number of bins in your table. You should start with a single bin, and then every time the total number of elements in the table divided by the number of bins is greater than K , you should double the number of bins and rehash the table. For uniformly distributed data, K will roughly correspond to the number of elements in each bin. Throw an `ece2400::InvalidArgument` if k is less than or equal to 0.

- `int Table<T,HashFunc>::size() const`
Return the current number of elements in the table. If the `Table<T,HashFunc>` is empty, this function should return 0.
- `void Table<T,HashFunc>::add(const T& value)`
Add the given value `value` to the `Table<T,HashFunc>`. You should first compute which bin the new value should be added based on applying the hash function to the given value. After determining the bin for the new value, call `Vector<T>::push_back` on the appropriate `Vector<T>` to add the value to the table. Then, check the load factor of your table. If necessary, create a new `Vector<Vector<T>>` with twice as many bins as the current one and rehash the values in the current `Vector<Vector<T>>` into the new one.
- `bool Table<T,HashFunc>::contains(const T& value) const`
Search the `Table<T,HashFunc>` for the given value and returns `true` if the value is found and `false` otherwise. If the `Table<T,HashFunc>` is empty, then this function should just return `false`. You should first calculate which bin will contain the given value, and then simply call `contains` on the corresponding `Vector<T>`.
- `Vector<T> Table<T,HashFunc>::to_vector() const`
Return a `Vector<T>` that contains all the items stored in the `Table<T,HashFunc>`. The `Vector<T>` should add all items from the first bin before adding the items from the second bin.
- `template <typename DistFunc>`
`const T& Table<T,HashFunc>::find_closest(const T& value, DistFunc dist) const`
Search the `Table<T,HashFunc>` and returns the value that has (approximately) the smallest difference from the given value. You should first calculate which bin will contain the given value, and then simply call `find_closest` on the corresponding `Vector<T>`. If the corresponding `Vector<T>` is empty, then you should return a default-constructed object of type `T`. If the `Table<T,HashFunc>` is empty, throw `ece2400::OutOfRange` with proper error message.
- `void Table<T,HashFunc>::print() const`
Print the content in the `Table<T,HashFunc>`. This member function is used for your own debugging purpose. You can implement this function in any way you like. You do not need to test this function. However, note that to make this member function generic you will need to use the C++ `iostream` library (e.g., `std::cout`) instead of `printf`. Consider printing the table such that it illustrates how the table is organized (i.e., print each bin a new line). We have provided you with a commented-out version of our implementation of this function that you can use to implement your own. You will need to replace the references to our private member fields with the names of your private member fields.

2.6. HRSLinearSearch Handwriting Recognition System

The first handwriting recognition system you will implement is `HRSLinearSearch`, which uses a brute force linear search algorithm. This system is almost identical to what you implemented in

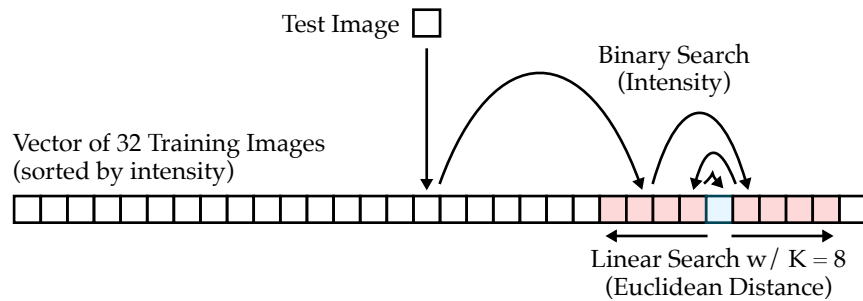


Figure 3: `VectorImage::find_closest_binary` **Example** – Example for 32 images and $K = 8$. Assumes vector images has already been sorted based on intensity. Binary search is based on intensity and quickly finds images with similar intensity as the test image. Linear search of 8 images is based on Euclidean distance and finds the closest match.

the previous programming assignment; the only difference is you should use `Vector<T>` instead of `VectorImage`. You are responsible for implementing each of the following functions:

```
HRSLinearSearch::HRSLinearSearch();
void HRSLinearSearch::train( const Vector<Image>& vec );
Image HRSLinearSearch::classify( const Image& Image );
```

Here is a brief specification for each member function of `HRSLinearSearch`:

- `HRSLinearSearch::HRSLinearSearch()`
Default constructor for `HRSLinearSearch`. Initialize your member variables if necessary.
- `void HRSLinearSearch::train(const Vector<Image>& vec)`
Train the HRS. For `HRSLinearSearch`, simply store a copy of the vector that contains the training images (`vec`). If you have implemented the assignment operator correctly, this should just be a one-line function.
- `Image HRSLinearSearch::classify(const Image& img)`
Classify the given image. This function should search through the entire training set and return the image that has the smallest euclidean distance from the given image. You should just call `Vector<Image>::find_closest_linear` with a function object that calculates the Euclidean distance between images.

2.7. `HRSBinarySearch` Handwriting Recognition System

The second handwriting recognition system you will implement is `HRSBinarySearch`, which uses a mix of binary and linear search. The overall approach is to first sort the training images by their intensity. Then during classification we can use binary search to quickly find training images with similar intensity, and then do a linear search of K images using Euclidean distance (see Figure 3). You are responsible for implementing each of the following functions:

```
HRSBinarySearch::HRSBinarySearch( int K = 1000 );
void HRSBinarySearch::train( const Vector<Image>& vec );
Image HRSBinarySearch::classify( const Image& Image );
```

Here is a brief specification for each member function of `HRSBinarySearch`:

- `HRSBinarySearch::HRSBinarySearch(int K)`
Default constructor for `HRSBinarySearch`. Initialize your member variables if necessary.
- `void HRSBinarySearch::train(const Vector<Image>& vec)`
Train the HRS. For `HRSBinarySearch`, you need to store a copy of the vector that contains the training images (`vec`), and then sort the training images based on intensity. You should just call `Vector<T>::sort` with a function object that compares the intensity of two images.
- `Image HRSBinarySearch::classify(const Image& img)`
Classify the given image (`img`). This function should use a binary and linear search. You should just call `Vector<T>::find_closest_binary` with a function object that calculates the Euclidean distance between images.

2.8. HRSTreeSearch Handwriting Recognition System

The binary search system requires sorting the entire array during training. We can potentially improve the performance of sorting by incrementally sorting the training dataset as we add each training image using a binary search tree. Given this intuition, the third handwriting recognition system you will implement is `HRSTreeSearch`, which should use the `Tree<Image, ImgCmpFunc>` data structure implemented earlier in the programming assignment.

You are responsible for implementing each of the following functions:

```
HRSTreeSearch::HRSTreeSearch( int K = 1000 );
void HRSTreeSearch::train( const Vector<Image>& vec );
Image HRSTreeSearch::classify( const Image& Image );
```

Here is a brief specification for each member function of `HRSTreeSearch`:

- `HRSTreeSearch::HRSTreeSearch(int K)`
Default constructor for `HRSTreeSearch`. Initialize the internal `Tree<Image, ImgCmpFunc>` data structure using the given `K` and a function that compares the intensity of images. We recommend declaring a functor as a private nested class inside `HRSTreeSearch`.
- `void HRSTreeSearch::train(const Vector<Image>& vec)`
Train the HRS. For `HRSTreeSearch`, simply iterate through all training images (`vec`) and add each one to the internal `Tree<Image, ImgCmpFunc>` data structure.
- `Image HRSTreeSearch::classify(const Image& img)`
Classify the given image. This function should simply call `Tree<Image, ImgCmpFunc>::find_closest` to return the image that has (approximately) smallest euclidean distance from the given image. You will need to pass in a function that calculates the Euclidean distance between images to `Tree<Image, ImgCmpFunc>::find_closest`.

Declaring a `Tree` type private data member in `HRSTreeSearch` requires the type of the image comparison function. We recommend using a functor declared as a private nested class inside `HRSTreeSearch` and declaring the data member like this:

```
Tree<Image, FunctorType> m_training_set;
```

2.9. HRSTableSearch Handwriting Recognition System

Both the binary search and tree search systems rely on searching a smaller subset of the entire training dataset to improve performance. We can also use a hash table to group similar images together, and then do a linear search through a smaller subset of similar images to improve performance. Given this intuition, the fourth handwriting recognition system you will implement is HRSTableSearch, which should use the Table<Image, ImgHashFunc> data structure implemented earlier in the programming assignment. You are responsible for implementing each of the following functions:

```
HRSTableSearch::HRSTableSearch( int K = 1000 );
void HRSTableSearch::train( const Vector<Image>& vec );
Image HRSTableSearch::classify( const Image& Image );
```

Here is a brief specification for each member function of HRSTableSearch:

- `HRSTableSearch::HRSTableSearch(int K)`
Default constructor for HRSTableSearch. Initialize the internal Table<Image, ImgHashFunc> data structure using the given K and a function that generates hash values for images. We recommend declaring a functor as a private nested class inside HRSTableSearch. The hash function you should implement is:
$$((\text{intensity} - 5,000) \% 20,000) * 100,000$$

This function distributes the images in our MNIST dataset into an approximately uniform distribution with values between 0 and INT_MAX. For more details on the hash function, see the end Of this section.
- `void HRSTableSearch::train(const Vector<Image>& vec)`
Train the HRS. For HRSTableSearch, simply iterate through all training images (vec) and add each one to the internal Table<Image, ImgHashFunc> data structure.
- `Image HRSTableSearch::classify(const Image& img)`
Classify the given image. This function should simply call Table<Image, ImgHashFunc>::find_closest to return the image that has (approximately) smallest euclidean distance from the given image. You will need to pass in a function that calculates the Euclidean distance between images to Table<Image, ImgHashFunc>::find_closest.

Declaring a Table type private data member in HRSTableSearch requires the type of the image hash function. We recommend using a functor declared as a private nested class inside HRSTableSearch and declaring the data member like this:

```
Table<Image, FunctorType> m_training_set;
```

Hash Function – Hash function design is an important part of developing an application using a table data structure. Different applications require different behaviors out of their hash functions. In our application, we want the bins in our table to store many images of similar intensity so that we can search through all of them for the closest match to the image we are trying to classify. We would also like to have the bins have an approximately equal number of items in them. This will make execution time consistent no matter what bin we are searching and ensure each classification looks at about the same number of images to find the closest match. Because we know something about the kind of data we will be storing in our table ahead of time, we can design a hash function to fit these criteria. The intensities of the images in the MNIST dataset are distributed like this:

```

4800 - 6400
6400 - 8000 #
8000 - 9600 ####
9600 - 11200 #####
11200 - 12800 #####
12800 - 14400 #####
14400 - 16000 #####
16000 - 17600 #####
17600 - 19200 #####
19200 - 20800 #####
20800 - 22400 #####
22400 - 24000 #####
24000 - 25600 #####
25600 - 27200 #####
27200 - 28800 #####
28800 - 30400 #####
30400 - 32000 #####
32000 - 33600 #####
33600 - 35200 #####
35200 - 36800 #####
36800 - 38400 #####
38400 - 40000 #####
40000 - 41600 #####
41600 - 43200 #####
43200 - 44800 #####
44800 - 46400 #####
46400 - 48000 ###
48000 - 49600 ##
49600 - 51200 ##
51200 - 52800 #
52800 - 54400 #
54400 - 56000

```

The minimum hash value is 5,086 and the maximum hash value is 79,483. These are not pictured on the histogram above because each symbol represents 100 intensity values in that range.

Our hash function tries to change this bell-shaped distribution with values between 5,000 and 80,000 and a peak at about 25,000 into an approximately uniform distribution with values between 0 and INT_MAX. The first thing our function does is to subtract 5,000 from the intensity. This will slide the bell curve towards 0, so that its values range from 0 to 75,000 with a peak around 20,000. Next, we use the remainder operator with a numerator of 20,000. This slides the upper half of the bell down to 0, and confines all values to the range 0 to 20,000. We now have a roughly normal distribution of values between 0 and 20,000, so to stretch it out and take full advantage of the range from 0 to INT_MAX we multiply by 100,000. The histogram of the hashes of all the values in the MNIST dataset using our hash function looks like this:

```

0 - 85899345 #####
85899345 - 171798690 #####
171798690 - 257698035 #####
257698035 - 343597380 #####
343597380 - 429496725 #####
429496725 - 515396070 #####
515396070 - 601295415 #####
601295415 - 687194760 #####
687194760 - 773094105 #####
773094105 - 858993450 #####
858993450 - 944892795 #####
944892795 - 1030792140 #####
1030792140 - 1116691485 #####
1116691485 - 1202590830 #####
1202590830 - 1288490175 #####
1288490175 - 1374389520 #####
1374389520 - 1460288865 #####
1460288865 - 1546188210 #####
1546188210 - 1632087555 #####
1632087555 - 1717986900 #####
1717986900 - 1803886245 #####
1803886245 - 1889785590 #####
1889785590 - 1975684935 #####
1975684935 - 2061584280 #####

```

You should feel free to experiment with your own hash function as a potential optimization to improve accuracy. If you do, make sure to be careful to avoid overflow and generating negative hash values!

2.10. HRSAlternative Handwriting Recognition System

For the final handwriting recognition system, students are free to experiment with any implementation they want. This alternative system needs to be significantly different from the four systems already explored in the programming assignments. For the alternative system only, students can use classes or functions from the standard C++ library. Students can also use third-party code that they find from other sources as long as they include proper credit both within the source code and their report. Note that the alternative system cannot rely on pre-compiled binaries. Any third-party code must be compiled as part of the project, which means students may need to make modifications to the included `CMakeLists.txt` file. The alternative system must be able to be built and tested both on `ecelinux`. It does not need to build on GitHub Actions. The alternative system can also rely on pre-computed training data. For example, students can use a special machine learning framework (e.g., PyTorch, TensorFlow) to pre-compute training information (e.g., neural network weights), store this training information in a file (e.g., text file, binary file) which is included in the Git repository, and then read from this file during the training phase using standard C++ file I/O.

Here are some ideas for possible alternative systems:

- **Parallel Linear Search** – Use the concurrent programming paradigm to have multiple threads searching different parts of the training dataset in parallel. You can start by hard coding the number of worker threads, and then later you can potentially add support for a configurable number of worker threads. Regardless, do not use more than three worker threads (i.e., four threads including the main thread) to avoid overloading the server. If you add support for a configurable number of worker threads, then you can experiment with how the performance scales as a func-

tion of the number of threads. Note that we do *not* recommend that students use a member function pointer as the callable object when initializing a thread. Member function pointers can be very challenging. We instead recommend using a free function, functor, or lambda. A lambda is probably the most elegant.

- **Parallel Linear Search with Thread Pool** – Use a “thread pool” instead of creating a new set of threads for every test image. A thread pool is a pool of persistent threads; create the threads in the constructor, then send commands to these “worker threads” to do searches, then do the join in the destructor. A thread pool will need to use synchronization primitives (e.g., atomics) to manage distributing work and determining when the worker threads are finished. Note that you can hard code the number of threads in your pool (i.e., 1–3) if you like, but again, do not use more than three worker threads (i.e., four threads including the main thread) to avoid overloading the server.
- **K Nearest Neighbors** – Instead of finding the closest match in the linear or binary search systems, find the K closest matches. Look at the labels of these K closest matches and vote. So if K is 5 and three or more of the five closest matches have the label 4, then classify the given test image as a 4. Experiment with optimizing how you store the current K closest neighbors to improve performance. Experiment with the trade-off in performance and accuracy as a function of K.
- **K-Means Clustering** – Use K-means clustering on the the training dataset to find K prototype images. Use majority voting within the k-th cluster to assign a label to each cluster (i.e., each prototype image). Note that K can be greater than 10. Compare the test image to the K prototype images and use the closest as a rough classifier. The same approach can be used to first compare to the K prototype images and then do a linear search through all images in that cluster to the test image.
- **MLP Neural Network** – Train a multi-layer perceptron neural network using a dedicated machine learning framework. Pre-compute the weights and bias terms, store them in a file, and then read this file into your system during the training phase. Use these weights in the classification phase.
- **Convolutional Neural Network** – Train a convolutional neural network using a dedicated machine learning framework. Pre-compute the weights and bias terms, store them in a file, and then read this file into your system during the training phase. Use these weights in the classification phase.
- **Support Vector Machine** – Train a support vector machine using a dedicated machine learning framework. Pre-compute all necessary information, store it in a file, and then read this file into your system during the training phase. Use this information in the classify phase.
- **Principle Components Analysis** – Vectorize an image into a 784 element vector and reduce its dimension via principle component analysis to say K dimensions. One can then use very simple classifier on the K dimensions, e.g., nearest neighbor search on the K-dimensional vector instead of the 784 dimensional one.
- **K-D Tree** – Use a k-d tree instead of a simple binary search tree to store the images. Tree each image as a 784-element vector. Randomly choose a dimension to use and then choose a threshold such that roughly halve the images are on each side of that threshold. Each leaf node should have M images, and you can use a linear search to find the closest match within these M images. Experiment with optimizing the trade-off in performance and accuracy as a function of M.
- **Locality-Sensitive Hashing** – Normally, hash tables choose hashing functions to avoid collisions. In locality-sensitive hashing we choose a hash function to encourage collisions for related items.

The table search is kind of like very simple locality-sensitive hashing. For “real” locality-sensitive hashing you will need to use a more sophisticated hash function. For example, you could take a small and random subset (but fixed random set) of pixels, quantize each pixel to 0 or 1 based on intensity, and then use these quantized pixels to create a binary number which you can then use as the index into your hash table. Experiment with different pixel subsets to find a good hash function which encourages collisions.

Note that almost all of these alternative systems will require students to leverage additional information external to the material covered in the course. This gives students an opportunity to connect the material in this course to what they might be learning in other courses, or to gain experience independently researching an algorithm and/or data structure and then putting it into practice. **For those students that wish to invest less time in implementing the alternative system, we recommend implementing the parallel linear search.** This specific alternative system requires applying what students learned in concurrent programming and should not require any additional information external to the material covered in the course. **Students can receive full credit for code functionality for a correct and efficient alternative system which uses concurrent programming!** If students are attempting an ambitious alternative system, they still might want to implement the parallel linear search first. This will enable students to have a “backup” in case they cannot get their more ambitious alternative system working.

Students are also free to submit multiple alternative systems. Simply create additional files named `src/HRSAAlternative2.h`, `src/HRSAAlternative2.cc`, `test/hrs-alternative2-directed-test.cc`, `eval/hrs-alternative2-eval.cc`, and update the `CMakeLists.txt` accordingly. **Note that we will only be evaluating the primary alternate system used in `eval/hrs-alternative-eval.cc`. So this should be your “best” version that you would the course staff to grade.**

3. Testing Strategy

You are responsible for developing an effective testing strategy to ensure all implementations are correct. Writing tests is one of the most important and challenging aspects of software programming. Software engineers often spend far more time implementing tests than they do implementing the actual program.

Note that while there are limitations on what you can use from the Standard C/C++ library in your *implementations* there are no limitations on what you can use from the Standard C/C++ library in your *testing strategy*. You should feel free to use the Standard C/C++ library in your golden reference models and/or for random testing.

3.1. Ad-hoc Testing

To help students start testing, we provide one ad-hoc test program per implementation in `src/sort-adhoc.cc`, `src/vector-adhoc.cc`, `src/image-adhoc.cc`, `src/tree-adhoc.cc`, and `src/table-adhoc.cc`. Students are encouraged to start compiling and running these ad-hoc test programs directly in the `src/` directory without using any build-automation tool (e.g., CMake and Make).

You can build and run the given ad-hoc test program for sort like this:

```
% cd ${HOME}/ece2400/netid/pa5-sys/src
% g++ -Wall -o sort-adhoc ece2400-stdlib.cc sort-adhoc.cc
% ./sort-adhoc
```

The `-Wall` flag will ensure that `g++` reports all warnings. For other ad-hoc tests you might need to include additional source files on the command line.

3.2. Systematic Unit Testing

While ad-hoc test programs help you quickly see results of your implementations, often too simple to cover most scenarios. We need a systematic and automatic unit testing strategy to hopefully test all possible scenarios efficiently.

In this course, we are using CMake/CTest as a build and test automation tool. For each implementation, we provide a directed test program that should include several test cases to target different categories and a random test program that should test that your implementation works for random inputs. **Unlike in the first three programming assignments, a great deal of tests have already been provided for you!** You can freely leverage the available tests to verify the functionality of your handwriting recognition systems. Remember however that your goal with respect to testing strategy is to convince yourself and the staff that your code is functional. If in order to convince yourself that your code is functional you realize further tests are needed (maybe just by copying and adjusting existing tests), then you should definitely write a few more.

We **strongly** encourage students to take an incremental design approach. **Do not implement all of these functions before running your first test!** Instead, we recommend students implement and test each of the following substeps.

- Step 1. Implement and test `sort<T, CmpFunc>`
- Step 2. Implement and test `Vector<T>`
 - Step 2a. Implement and test the default constructor, destructor, `push_back`, `size`, `at`
 - Step 2b. Implement and test the non-default constructor
 - Step 2c. Implement and test `contains`
 - Step 2e. Implement and test `sort`
 - Step 2d. Implement and test `find_closest_linear`
 - Step 2f. Implement and test `find_closest_binary`
 - Step 2g. Implement and test `operator []`
 - Step 2h. Implement and test the copy constructor
 - Step 2i. Implement and test the assignment operator
- Step 3. Implement and test `Image`
 - Step 3a. Implement and test the default/non-default constructor, `get_ncols`, `get_ncols`, `at`
 - Step 3b. Implement and test `operator []`
 - Step 3c. Implement and test `set_label`, `get_label`
 - Step 3d. Implement and test `get_intensity`
 - Step 3e. Implement and test `operator ==`, `operator !=`
- Step 4. Implement and test `Tree<T, CmpFunc>`
 - Step 4a. Implement and test the default constructor, destructor, `add`, `size`, `contains`
 - Step 4b. Implement and test `to_vector`
 - Step 4c. Implement and test `find_closest`
 - Step 4d. Implement and test the copy constructor
 - Step 4e. Implement and test the assignment operator
- Step 5. Implement and test `Table<T, HashFunc>`
 - Step 5a. Implement and test the default constructor, destructor, `add`, `size`, `contains`
 - Step 5b. Implement and test `to_vector`

- Step 5c. Implement and test `find_closest`
- Step 6. Implement and test `HRSLinearSearch`
- Step 7. Implement and test `HRSBinarySearch`
- Step 8. Implement and test `HRSTreeSearch`
- Step 9. Implement and test `HRSTableSearch`
- Step 10. Implement and test `HRSAlternative`

Each substep should correspond to a directed test case. Note that for some of these steps we provide both *generic* and *specialized* tests. For example, if you look in `sort-directed-test.h` you will see generic tests which are templated on the type `T` and the comparison function `CmpFunc`. This enables us to reuse the exact same test functions to test sorting an array of ints and to test sorting an array of Images. You will want build and run the specialized test programs (i.e., `sort-int-directed-tests` and `sort-image-directed-tests`).

As in the previous programming assignment, we provide you a testing framework you should use for your directed and random testing. See the provided test programs in the `test` subdirectory for how to use this framework. The ECE 2400 standard library in `ece2400-stdlib.h` contains the following macros you should use to check the correctness of your implementations:

- `ECE2400_CHECK_FAIL()` – check program does not reach this point
- `ECE2400_CHECK_TRUE(expr_)` – check `expr_` is always true
- `ECE2400_CHECK_FALSE(expr_)` – check `expr_` is always false
- `ECE2400_CHECK_INT_EQ(expr0_, expr1_)` – check `expr0_ equals expr1_`

Before running the tests you need to create a separate build directory and use `cmake` to create the Makefile like this:

```
% cd ${HOME}/ece2400/groupid/pa5-sys
% mkdir -p build
% cd build
% cmake ..
```

Now you can build and run all unit tests for all implementations like this:

```
% cd ${HOME}/ece2400/groupid/pa5-sys/build
% make check
```

If you are failing a test program, then you can “zoom in” and run all of the unit tests for a single test program (e.g., directed tests for `sort` specialized for `int`) like this:

```
% cd ${HOME}/ece2400/groupid/pa5-sys/build
% make sort-int-directed-test
% ./sort-int-directed-test
```

You can then “zoom in” to a specific test case by passing in the index of that test case like this:

```
% cd ${HOME}/ece2400/groupid/pa5-sys/build
% make sort-int-directed-test
% ./sort-int-directed-test 1
% ./sort-int-directed-test 2
```


Your implementation of `find_closest` for both `Tree<T,CmpFunc>` and `Table<T,HashFunc>` may be slightly different than the staff implementation. These differences may not significantly impact the overall accuracy of your handwriting recognition systems, but might cause you to fail one or two of the test assertions in the directed testing. In this case, it is acceptable for you to carefully update the directed testing to better match your specific implementation. You might also consider adding some random testing for `find_closest` for both `Tree<T,CmpFunc>` and `Table<T,HashFunc>`.

3.3. Memory Leaks

Students are also responsible for making sure that their program contains no memory leaks or other issues with dynamic allocation. We have included a make target called `memcheck` which runs all of the test programs with `Valgrind`. `Valgrind` will force the test to fail if it detects any kind of memory leak or other issues with dynamic allocation.

You can check memory leaks and other issues with dynamic memory allocation for all your test programs like this:

```
% cd ${HOME}/ece2400/groupid/pa5-sys/build
% make memcheck
```

You can just check one test program (e.g. `sort-int-directed-test`) like this:

```
% cd ${HOME}/ece2400/groupid/pa5-sys/build
% make sort-int-directed-test
% valgrind --trace-children=yes --leak-check=full \
  --error-exitcode=1 --undef-value-errors=no ./sort-int-directed-test
```

Those are quite a few command line options to `Valgrind`, so we have created a ‘`ece2400-valgrid`’ script. This script is just a simple wrapper which calls `Valgrind` with the right options.

```
% cd ${HOME}/ece2400/groupid/pa5-sys/build
% make sort-int-directed-test
% ece2400-valgrid ./sort-int-directed-test
```

3.4. Code Coverage

After your implementations pass all unit tests, you can evaluate how effective your test suite is by measuring its code coverage. The code coverage will tell you how much of your source code your test suite executed during your unit testing. The higher the code coverage is, the less likely some bugs have not been detected. You can run the code coverage like this:

```
% cd ${HOME}/ece2400/groupid/pa5-sys
% rm -rf build-coverage
% mkdir -p build-coverage
% cd build-coverage
% cmake ..
% make check
% make coverage
```

Note that these code coverage results will reflect *all* prior runs of the test and evaluation programs in the build directory. That is why in the above example, we do a fresh build in a separate `build-coverage` build directory.

If you want to drill down and explore the coverage of each line in a program you use use the elinks web browser like this:

```
% cd ${HOME}/ece2400/groupid/pa5-sys/build-coverage
% elinks coverage-html/index.html
```

Code coverage is just one more piece of evidence you can use to make a compelling case for the correct functionality of your implementations. It is not required that students achieve 100% code coverage. It is far more important that students simply use code coverage as a way to guide their test-driven design than to overly focus on the specific code coverage number.

4. Evaluation

In the first three PAs, students gained first-hand experience with test-driven design through implementing a compelling testing strategy. The final two PAs shift to focus more on incremental profiling and optimization.

Once you have verified the functionality of your handwriting recognition systems, you can evaluate their performance and accuracy with breakdowns for both the training phase and the classification phase. You can build the evaluation programs like this:

```
% cd ${HOME}/ece2400/groupid
% mkdir -p pa5-sys/build-eval
% cd pa5-sys/build-eval
% cmake -DCMAKE_BUILD_TYPE=eval ..
% make eval
```

Note that we are working in a separate `build-eval` directory, and that we are using the `-DCMAKE_BUILD_TYPE=eval` command line option to the `cmake` script to create optimized executables without any extra debugging information. **You must do your quantitative evaluation using an eval build. Using a debug build for evaluation produces meaningless results.**

The following runs a complete evaluation with 60K training images, 10K testing images, and K set to 1000 for each of the five handwriting recognition systems.

```
% cd ${HOME}/ece2400/groupid/pa5-sys/build-eval
% make eval
% ./hrs-linear-search-eval
% ./hrs-binary-search-eval
% ./hrs-tree-search-eval
% ./hrs-table-search-eval
% ./hrs-alternative-eval
```

You can also specify the size of the number of training images, the number of classification images, and the size of the final linear search when using binary search on the command line for each evaluation program:

```
% ./hrs-linear-search-eval N M
% ./hrs-binary-search-eval N M K
% ./hrs-tree-search-eval N M K
% ./hrs-table-search-eval N M K
```

```
% ./hrs-alternative-eval N M
```

Where N is the number of training images, M is the number of classification images, and K is the size of the final linear search when using the binary search. Note that, using a smaller number of training images and/or classification images is only for profiling and iterative performance optimization. All accuracy results must use the full 60K training dataset and 10K classification dataset.

You should leverage your insights from the previous programming assignment to apply similar optimizations that you feel are worthwhile to these handwriting recognition systems. You may need to do some performance profiling to figure out where the bottleneck is and what to optimize. You can use `perf` to create a flame graph for each implementation. *You should only use `perf` if the execution time is about five minutes or less. If you use `perf` when the execution time is longer it will create a huge trace file which will fill up your home directory!* You can create the flame graph for the `HRSLinearSearch` with the following command:

```
% cd ${HOME}/ece2400/groupid/pa4-sys/build-eval
% perf record --call-graph dwarf ./hrs-linear-search-eval
% perf script report stackcollapse | flamegraph.pl > graph-linear.svg
```

While you are encouraged to use flame graphs to guide your optimization, you do not need to include them in your report unless you think it would be useful in your quantitative evaluation section. Unlike the previous programming assignment, you do not need to record incremental performance after every optimization. You should focus on analyzing the final optimized performance results in your report. You should optimize the linear and binary search systems so that they can achieve the performance targets used in the previous programming assignment.

Note that it is perfectly fine if your alternate system is slower and less accurate compared to the linear, binary, tree, and table search systems. You can still receive *full* credit for code functionality *and* full credit for optimization, even if your alternate system is slower and less accurate compared to the other systems. The goal of the alternate system is to have fun and try something different. It can be hard to make these alternate systems significantly faster and more accurate. You might need to adjust the directed tests to have a lower accuracy threshold for your alternate system, which is acceptable as long as you can justify why the accuracy is lower. We will award a small bonus for the best alternate systems, but again, you can receive full credit for functionality and optimization even if your alternate system turns out to be slower and less accurate. Of course you should discuss what you implemented in the alternative implementation section of your report and *why* your alternate system did or did not do as well as you hoped in the quantitative evaluation section of your report.

5. Putting It All Together

Recall that complete systems often include a frontend written in a productivity-level language (e.g., Python) and a backend written in an efficiency-level language (e.g., C/C++). We have provided you a Python GUI frontend which allows you to classify real handwritten digits using one of the five available backends (`HRSLinearSearch`, `HRSTreeSearch`, `HRSTableSearch`, and `HRSAAlternative`). To use this Python GUI frontend, first make sure that you have built the `hrs-backend` target in the build system (preferably an evaluation build since it will be faster!):

```
% cd ${HOME}/ece2400/groupid/pa5-sys
% mkdir -p build-eval
% cd build-eval
% cmake .. -DCMAKE_BUILD_TYPE=eval
```

```
% make hrs-backend
```

Copy the compiled backend program to the evaluation directory:

```
% cd ${HOME}/ece2400/groupid/pa5-sys
% cp build-eval/hrs-backend eval/hrs-backend
```

Then launch the GUI frontend (hrs-frontend.py) using Python:

```
% cd ${HOME}/ece2400/groupid/pa5-sys/eval
% python hrs-frontend.py
```

The frontend only works if you use a remote access option that supports Linux applications with a GUI. If you are using VS Code and you want to experiment with the GUI, then you will need to use either X2Go, MobaXterm, or Mac Terminal with XQuartz. See the course tutorials for more information on these alternative remote access options.

6. Milestone and Report

This section includes critical information about the incremental milestone, final code submission, and the final report specific to this PA. **The programming assignment logistics document provides general details about the requirements for the milestone and final submission.** You must actually read the document to ensure you know how we will access your milestone and final submission.

6.1. Incremental Milestone

While the final code and report are all due at the end of the assignment, we also require you to complete an incremental milestone, push and submit your code to GitHub on the date specified by the instructor. More specifically to meet the incremental milestone of this PA, you are expected to:

- Complete the implementation of `sort<T, CmpFunc>`
- Complete the implementation of `Vector<T>`
- Complete the implementation of `Image`
- Pass all given directed and random tests for these implementations
- Consider adding a few of your own directed tests

Here is how we will be testing your milestone:

```
% mkdir -p ${HOME}/ece2400/submissions
% cd ${HOME}/ece2400/submissions
% rm -rf repo
% git clone git@github.com:cornell-ece2400/groupid

% cd ${HOME}/ece2400/groupid/pa5-sys
% mkdir -p build
% cd build
% cmake ..
% make check-milestone
```

6.2. Final Code Submission

Your code quality score will be based on the way you format the text in your source files, proper use of comments, deletion of instructor comments, and uploading the correct files to GitHub (only source files should be uploaded, no generated build files). To assist you in formatting your code correctly, we have created a make target that will autoformat the code for you. You can use it like this:

```
% cd ${HOME}/ece2400/groupid/pa5-sys
% mkdir -p build
% cd build
% cmake ..
% make autoformat
% git diff
# ... check all changes ...
% git commit -a -m "autoformat"
```

Note that the autoformat target will only work if you have already committed all of your work. This way you can easily use `git diff` to view the changes made by the autoformatting and commit those changes when you are happy with them. Since we provide students an automated way to format their code correctly, students have no excuse for not following the course coding conventions!

Note that students must remove unnecessary comments that are provided by instructors in the code distributed to students. Students must not commit executable binaries or any other unnecessary files. The autoformat target will not take care of these issues for you.

To submit your code you simply upload it to GitHub. Your code will be assessed both in terms of functionality and code quality. Your functionality score will be determined by running your code against a series of tests developed by the instructors to test its correctness. Here is how we will be testing your final code submission:

```
% mkdir -p ${HOME}/ece2400/submissions
% cd ${HOME}/ece2400/submissions
% rm -rf repo
% git clone git@github.com:cornell-ece2400/netid

% cd ${HOME}/ece2400/submissions/groupid/pa5-sys
% mkdir -p build
% cd build
% cmake ..
% make check
% make memcheck
% make eval
# ... run the eval programs ...
```

6.3. Final Report

The final report must be uploaded to Canvas. The date you upload your report will indicate how many slip days you are using for the assignment. For this PA, we require you to include five sections: introduction, alternative implementation, complexity analysis, quantitative evaluation, conclusion, and work distribution. There is no page limit on this report.

		HRSTreeSearch	HRSTableSearch
training	execution time	$T_K(N)$	
training	time complexity w.r.t. N	big-O	
classify one image	execution time	$T_K(N)$	
classify one image	time complexity w.r.t. N	big-O	
training + classify M images	execution time	$T_K(N, M)$	
training + classify M images	time complexity w.r.t. N	big-O	
training + classify M images	time complexity w.r.t. M	big-O	

Table 1: Template for Complexity Analysis Table

The alternative implementation section should describe at a high-level your alternative handwriting recognition system. Think critically about what are the key items to mention in order for the reader to understand how the alternative implementation works. Motivate why you decided to pursue this specific alternative implementation. Examples are usually great to include here to illustrate how the alternative implementation works. Students are highly encouraged to include pseudo-code where appropriate. Do not include C++ code; your report should be at a higher level. Students must provide a balanced discussion of not just the implementation itself, but why you chose to take this approach.

The complexity analysis section of your report must include a table that includes the execution time of the training and classification phases as a function of N and K . The table should include the time complexity (in big-O notation) for the training and classification phases with N as the input variable. The table should include the execution time for training and classifying with N training images and M test images as a function of N , M , and K . Finally, the table should include the time complexity (in big-O notation) when training and classifying with N training images and M test images. In this situation, we have two input variables, so you should first analyze the time complexity with respect to N assuming M is a constant (i.e., we will classify a fixed number of images while the number of training images grows large), and then analyze the time complexity with respect to M assuming N is a constant (i.e., we have a fixed number of training images as the number of images we wish to classify grows large). See Table 1 for a template. Justify your table entries. You should discuss (at a high level) your chosen units. We recommend using number of images accessed in the classification dataset. Note that you don't need to explicitly discuss all six steps of complexity analysis and we are not looking for a rigorously formal proof, but you do need to be clear about the assumptions you made during analysis and provide some kind of compelling high-level argument.

The quantitative evaluation section of your report should include the following table and plot.

- *Experimental Table of Execution Time and Accuracy*: Table with one column for (optimized) training time, (optimized) classification time, (optimized) total execution time, and accuracy and one row for each of implementation. Use the full ($N = 60K$), full classification dataset ($M = 10K$), and set K to 1000.

- *Experimental Plot of Execution Time vs. Accuracy:* Plot with total execution time of the training plus classification phases on the y-axis vs. accuracy on the x-axis for the full ($N = 60K$) and full classification dataset ($M = 10K$). The plot should have one point for `HRSLinearSearch` and five points for `HRSBinarySearch`, `HRSTreeSearch`, and `HRSTableSearch` each for a different K . Make sure you label the points with the corresponding value of K . We recommend using $K = 1, 10, 100, 1000, 10000$. The plot should also have one or more points for the alternate system, since your alternate system might also include a parameter that can be used to trade-off accuracy and performance. For example, if you chose to implement a parallel linear search for your alternative system, then you might consider including three points in your plot (i.e., one worker thread, two worker threads, and three worker threads).

The quantitative evaluation section should be structured into the following paragraphs (they don't need to be super long paragraphs!).

- *Paragraph 1:* Discuss the performance of `HRSBinarySearch`. Use the performance table to discuss the trade-off between time spent during training vs. time spent during classification compared to `HRSLinearSearch`. Discuss how the performance and accuracy vary as a function of K . Connect these quantitative results to your complexity analysis from the previous programming assignment. Explain *why* performance and accuracy vary as a function of K . Feel free to bring in some of your analysis from the previous programming assignment.
- *Paragraph 2:* Discuss the performance of `HRSTreeSearch`. Use the performance table to discuss the trade-off between time spent during training vs. time spent during classification compared to `HRSLinearSearch`. Discuss how the performance and accuracy vary as a function of K . Connect these quantitative results to your complexity analysis section. Explain *why* performance and accuracy vary as a function of K .
- *Paragraph 3:* Discuss the performance of `HRSTableSearch`. Use the performance table to discuss the trade-off between time spent during training vs. time spent during classification compared to `HRSLinearSearch`. Discuss how the performance and accuracy vary as a function of K . Connect these quantitative results to your complexity analysis section. Explain *why* performance and accuracy vary as a function of K .
- *Paragraph 4:* Compare and contrast all handwriting recognition systems except for `HRSAternative`. Explicitly discuss the trade-off between accuracy and performance across all handwriting recognition systems. Explain *why* one handwriting recognition system performs better in terms of execution time and/or accuracy. Consider including the notion of *pareto-optimal* in your analysis.
- *Paragraph 5:* Discuss the performance of `HRSAternative`. Use the performance table to discuss the trade-off between time spent during training vs. time spent during classification compared to `HRSLinearSearch`. Compare and contrast `HRSAternative` to the other handwriting recognition systems.
- *Paragraph 6:* Discuss a specific application where we might want to use a handwriting recognition system. What platform (i.e., cloud server, laptop, smart phone, internet-of-things device, etc.) would we use to run this application? Why would this application need handwriting recognition? What is a reasonable execution time and accuracy target for the application in this context? Does one or more of the handwriting recognition systems implemented in this programming assignment meet these targets? Keep in mind that you need to divide the classification time by $M = 10K$ to find the classification time per image. Be quantitative and specific. **(This is a very important paragraph!)**

The work distribution section should be a one paragraph description of which student did what work. It is perfectly fine if one students does more work; the key is to transparent and honest about the work distribution across students.

Acknowledgments

This programming assignment was created by Christopher Batten, Christopher Torng, Tuan Ta, Yanghui Ou, Peitian Pan, and Nick Cebry as part of the ECE 2400 Computer Systems Programming course at Cornell University. We also thank the curators of the MNIST database of handwritten digits.