

Communicating Mathematics via Pen-Based Interfaces

Elena Smirnova⁺ and Stephen M. Watt*
Ontario Research Centre for Computer Algebra
The University of Western Ontario
London Ontario, Canada N6A 5B7
watt@csd.uwo.ca

Abstract—We address the question of how to organize pen-based interfaces for mathematical software systems, both for computational mathematics packages and document processing software. Our architecture includes components for ink collection, mathematically-oriented recognizers, portability support and interfaces to applications. We summarize aspects of mathematical handwriting recognition and discuss the methods we have used for individual character recognition and overall expression analysis. We present our pen-based computing environment, Mathink, and give an overview of facilities for training, ink annotation, and testing.

I. INTRODUCTION

It is natural to ask what might be the most effective computer interface to handle mathematics, be it for a symbolic computation system or a word processor. This question has been broached by various authors, including [33], [34], and one of the conclusions is that many users would prefer to enter expressions using a pen, rather than a keyboard. Although there has been a long-standing interest in this problem (*e.g.* [2], [14]), the recent proliferation of pen-enabled devices, such as personal digital assistants, tablet PCs and interactive whiteboards, all supported by the processing power of modern computers, suggests that pen-based support for mathematics could now reach a wide audience. From a certain point of view, the digital tablet is favorable for mathematical input compared to conventional tools such as paper or chalkboard, owing to the rich variety of software possible with ink-capturing hardware. Mathematical content entered with a digital pen can be processed in many interesting ways, including by editing, validation and semantically-driven direct manipulation, as well as by interactive exploration of alternative hypotheses, recording derivations, and so on.

These observations imply that pen-based interfaces for mathematics must have certain capabilities, including collecting and processing of digital ink, recognition of handwritten expressions and connection to mathematical engines. Furthermore, pen-based computing occurs in a context that is rich with standards for representing and communicating mathematical data, including Unicode [27], MathML [1] and InkML [5]. We

see support of these standards as important for cross-platform communication and collaboration.

This paper explores the range of issues that arise in constructing such pen-based interfaces for mathematics. Our presentation is organized as follows: Section II presents our architecture for mathematical pen interfaces. Section III describes the main features of a pen-based mathematical computing system, Mathink, that we have developed to evaluate our architectural approach. Section IV discusses aspects of the recognition process specific to mathematical handwriting. Next, in Section V, we consider two mechanisms for communicating between pen-based front-ends and mathematical computation engines. Finally, we describe the tools we have developed for system training and testing, and show how they can be used to improve performance in recognition.

II. A PORTABLE FRAMEWORK FOR PEN COMPUTING

We begin by summarizing the principal purposes and key requirements of systems built for pen-based computing. We describe a framework architecture designed to implement such systems allowing high-quality handling of digital ink, while ensuring portability across platforms and applications.

A. Objectives

The objective of our work has been to identify and investigate the issues whose resolution will lead to effective pen-based mathematical computation. This has meant investigating a number of questions and approaches to different problems, from ink collection, to character recognition, expression analysis and manipulation. From the outset, we have recognized that in each of these areas there are important, difficult questions and that a full solution will require work in each of these areas. Consequently, the software system Mathink (described later) that we have developed during the course of this study has been intended as a testbed to explore ideas, rather than as a production environment.

Our goal was not to create a stand-alone mathematical recognizer nor to develop an application-specific plug-in. On the contrary, we have studied how to provide a uniform interface to many applications through one component. For concreteness, we have identified two classes of hosting environments: computer algebra systems and rich document editors.

⁺ Now at Texas Instruments, e-smirnova@ti.com.

* Supported in part by the Natural Sciences and Engineering Research Council of Canada, Microsoft Canada and Maplesoft.

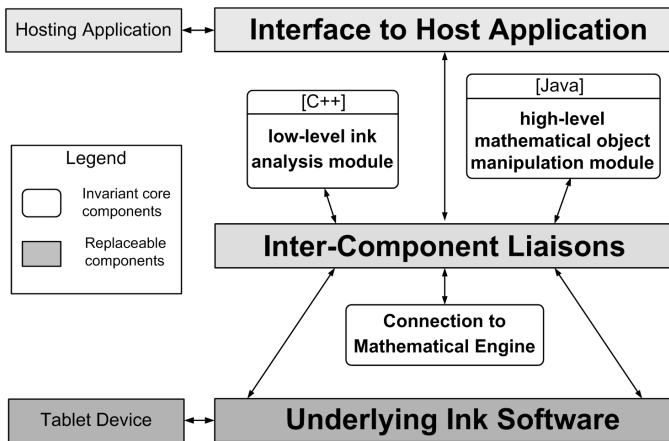


Fig. 1. Top-level framework organization

Modeling a mathematical ink-handling component as one that may be used in many contexts places a number of constraints on the system architecture design. We require that such a pen-based interface be portable across a range of platforms without sacrificing digital ink quality. In previous work [21], we have presented an architectural solution that allows one to embed device-specific support for digital ink and to ensure compatibility with hosting applications on a variety of platforms. What we present here is based on this approach.

B. Architecture Organization

The key point of our design is to separate the components responsible for the analysis of handwritten mathematics from the modules that provide connection to hosting applications and underlying platforms, as shown in Figure 1. In this way, core recognition components can remain invariant, while interfaces and adaptors can be replaced for each combination of platform and hosting application. This allows re-use of the main capabilities of the pen-based framework in different environments without changing their internal organization.

With this approach, development and maintenance of recognition modules requires expertise only in mathematical handwriting analysis. On the other hand, implementation of the “glue” components to integrate recognition units into hosting environments does not demand any specific knowledge in mathematical recognition. Therefore docking mechanisms can be developed independently by systems experts.

In our experiments with target platforms and applications, we have arrived at two core modules for character recognition and structural analysis. Our choice of implementation languages for these components had to satisfy two principal criteria: platform portability and support of the functionality required by the modules. For the character recognition unit, we chose C++ because of the quality of compilers for these computationally intensive tasks. Structural analysis, however, involves less computation and benefits from rich functionality for expression manipulation and XML support; therefore, for the second module we chose to use Java.

Another important component of the framework is an interface to mathematical engines. As well as recognizing ink input and translating it into a valid mathematical expressions, we expect our pen-based component to allow further computation and manipulation with the formulæ recognized. To provide this functionality we need to allow a back-end for mathematical computing in our framework. Instead of building yet another symbolic computation system we have decided to provide access to already existing powerful and well-developed symbolic computing packages. To satisfy the portability criteria, among the available computer algebra systems we have considered only those which are supported on multiple platforms. We discuss connection mechanisms to mathematical mathematical software systems in Section V.

III. THE MATHINK SYSTEM

We now introduce *Mathink*, an ink-aware mathematical component implemented as a test of our architectural approach. *Mathink* is designed both as mathematical pen-based plug-in and as an experimental environment to train and test mathematical recognizers. We briefly review the top-level interface and present the connection mechanisms that allow *Mathink* to be used as a control running inside MAPLE [13] worksheets and Microsoft Word documents.

A. *Mathink* from a Bird’s Eye Perspective

Mathink provides a user interface allowing digital ink to be collected from a pen, mouse or other device. The system supports a range of ink source types, including both pressure-sensitive and touch-activated digitizers. We have successfully tested the system with Tablet PCs, SmartBoards, the e^3 works Stylo USB tablet and Wacom digitizers. This design allows the system to be adapted also for PDAs, since they support ink collecting protocols similar to those used by SmartBoard. In a Microsoft environment, the system uses the Tablet PC SDK to collect high-resolution digital ink.

The flow of control through the *Mathink* system is organized as follows: Ink glyphs are entered in the writing area is transferred to the system after an adjustable time delay. The writer thus is provided with immediate feedback from the recognizer for every character entered. Along with the best match, the recognizer offers other high ranked candidates. These are shown in a fixed location that allows rapid selection of alternatives, whose total number can be preset by the user.

Recognition results accepted by the user are displayed super imposed on or replacing the original ink. Having these two options allows users to switch between showing the ink input and typeset recognition results. When in typeset mode, the recognized content is drawn within the bounding box of the original ink strokes. The best fit is calculated by the *Mathink* rendering module. This involves adjusting the size of the characters on a given baseline, while taking into account positioning of large grouping operators, fractions and radicals.

The layout of an expression is re-analyzed with each new character entered. Once the user has finished writing a formula, the system finalizes and refines its structure. It is then

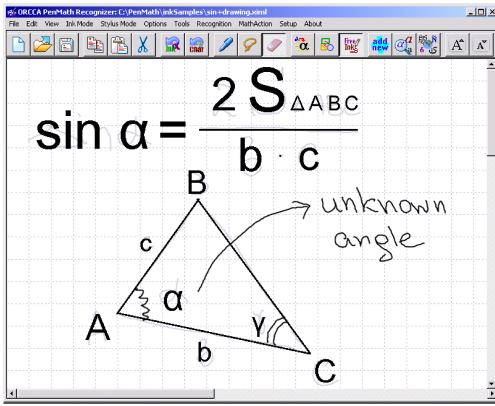


Fig. 2. Mathematical ink document with multiple recognition options

parsed to a standard mathematical format, such as Presentation MathML [1]. Once the mathematical content is constructed from handwritten input, the user can choose to send it to a symbolic computation back-end to perform further operations such as evaluation, simplification, factorization, solving, *etc.*

In addition to its main purpose of manipulating mathematical expressions, the Mathink system allows the user to make sketches or diagrams. In drawing mode ink input is recognized as basic geometrical shapes instead of mathematical characters. In addition, Mathink also allows to enter “free ink” that is not sent to the recognizer. Switching between the modes is permitted at any time, so a combination of recognized formulae, geometrical drawing, quick notes and arbitrary sketches can be placed in the same document, as shown in Figure 2.

As opposed to the approach taken in the Infty Editor [7], we do not discard collected ink after it has been recognized. The original strokes along with their annotations can be saved, either in Microsoft Ink Serialized Format (ISF) or in the portable system-independent standard format InkML [5]. Stored ink can be re-opened later in the Mathink editor or played back using the InkPlayer tool. This possibility allows the evaluation of different combinations of recognizers and settings on the same hand-written examples. In particular, we use this ability to test and optimize the performance of the system, as discussed in Section VII.

B. Exporting Mathink as a Plug-In Control

As we stated at the beginning, our main goal has been to enable pen-based mathematical interfaces to computer algebra systems, rich document editors and other applications. We explored methods for the integration of ink systems with various software packages. Among the possible approaches, we studied methods to export the Mathink interface as a standard plug-in control compatible with target applications.

To do this, we defined a subset of the Mathink functionality to expose to the hosting systems. These covered (a) manipulation of digital ink, including collection, rendering and clipboard operations, (b) tuning recognizer settings, and (c) exporting recognition results both for individual glyphs

and for whole expressions. The rest of the interface functions were designed to manage Mathink control properties, such as appearance and accessibility.

While the functionality exposed by a plug-in control should remain as invariant as possible across applications, the type and internal organization of a plug-in component in most cases is determined by the hosting system. We chose Microsoft Office as a suitable platform to evaluate our approach with document processors. As the second independent approach, we experimented with MAPLE 10 to explore a pen-based mathematical interface with a computer algebra package. These choices led us to two forms of plug-ins: as an *ActiveX control* to be used with Microsoft applications and with *JavaBeans* to connect the Mathink interface in MAPLE.

Presenting the Mathink component as an *ActiveX control* involved developing a number of interacting adaptors, implemented using Microsoft “COM Interop” [6] and “Platform Invoke Services” [16]. Since .NET managed code cannot be directly compiled into an ActiveX control, we had to add an intermediate *Win32* component to host the .NET ink collector. The installer for the Mathink plug-in on the Windows platform integrates it into Microsoft Word or Excel as an “ORCCA Mathink Control” and places a new button on the “Formatting” toolbar. In other Microsoft applications, the Mathink interface is accessible as a COM object, for example, through the “Insert|Object..” menu item.

Providing the Mathink interface as a *JavaBean* required connection between the .NET and Java platforms. Existing commercial and open source packages providing a link between .NET and Java are known to lag behind one or both of the platforms as they evolve. Our approach uses no third party software. The connection mechanisms we have developed are based on standard protocols that are well-defined and properly maintained for both platforms. While remaining simple, this solution suits our purposes. Because the Mathink system involves a fair amount of managed .NET code, first we had to expose it via COM Interop as a COM object. The “unmanaged” nature of COM components allows their exposure to Java via the *Java Native Interface (JNI)*. We use the *Abstract Window Toolkit Native Interface* [3] to permit rendering to Java canvas from native code. Finally, we wrapped a JNI adaptor within a Java package and exported its main classes as *JavaBeans*. Once the Mathink control is enabled on a Java platform, it can be incorporated into MAPLE worksheets. This not only allows the Mathink component to serve as a pen interface to the MAPLE computer algebra system, but also provides a direct connection to the MAPLE kernel, as discussed in Section V.

IV. MATHEMATICAL HANDWRITING RECOGNITION

The past four decades have seen mathematical handwriting recognition rise as an area of research interest. Recent results on aspects of mathematical handwriting analysis and recognition have been reported by various groups [7], [15], [17], [18], [32], including our own [4], [8], [9], [10], [11], [23], [24], [28], [29], [30], [31] and our colleagues at the University of Waterloo [12], [19].

Handwriting analysis for mathematical content is commonly divided into four interacting stages: collection of digital ink, recognition of individual characters, structure analysis, and determination of mathematical semantics. In this section we describe the techniques used we have used to translate ink traces into meaningful mathematical formulæ to be used by computer algebra packages, typesetting systems or equation editors within document processing environments.

A. Collecting Ink

Our system assumes that every character is recognized immediately after it has been written. This ensures that when the user continues writing, all previous characters have been recognized correctly. A disadvantage of this approach is that we partially give up a pure “pen and paper” paradigm, because the user has to wait for each character to be recognized. This style of entering ink also has benefits, however. It avoids the user having to search through lengthy formulæ for mis-recognized entries after large expressions are processed. On the implementation side, entering one character at a time allows the procedure of grouping strokes into glyphs to be skipped. Stroke grouping otherwise presents a complex problem. Another advantage of the “one character at a time” design for mathematical recognizers is that it is possible to make use of a context information, as shown in [24]. Mathematical context is derived from the previously entered characters. Therefore, if the latest recognition results is not confirmed by the writer, then processing the context becomes much more complex.

B. Character Recognition

After collecting the ink for a character, we extract the trace data and pass it to the recognizer engine. The recognition has three phases: First, the raw ink is preprocessed. Then the refined ink is analyzed to filter out obvious mismatches from the prototype set. In the final stage, handwritten input is compared with a set of known ink models, and the best matches are returned as recognition candidates. We describe each of these steps in more details below.

1) *Ink Preprocessing*: Preprocessing first eliminates noise that is often present at the beginning and at the end of strokes. Next, ink glyphs are normalized by resampling and resizing. *Resampling* helps to uniformize distances between points in the trace. Point distributions in strokes are not consistent due to a variance in sampling rates on different ink devices. Variations in users’ writing speeds also affect the density of trace points. Resampling reduces the amount of data that must be retained without reducing the fidelity and makes the later matching more uniform. *Trace resizing* assures that the size of a written ink glyph does not affect the recognition results. The final stage of ink preprocessing is to smooth strokes to eliminate any accidental artefact produced during cleaning and normalization.

The most recent version of our preprocessor also includes options for stroke *re-direction* and *reordering* to allow the recognizer to match characters whose shapes are identical,

but writing styles are different. With this feature enabled, the letter *X* would be recognized correctly regardless of which stroke was entered first and whether the strokes were written downwards or upwards.

2) *Feature analysis*: This step extracts features of the written glyphs. These features are used to categorize ink by overall appearance, geometrical shape and writing style. Features related to appearance include, for example, the height to width ratio; geometrical properties are based on the number of cusps, loops and intersection. Features related to writing style include number of strokes, directions and spatial point density. These criteria are used for prototype pruning, which breaks a large model dataset into smaller classes, each containing only 12-15% of all models. This consequently increases recognition performance. In [31] and [30] Watt and Xie describe feature extraction and prototype pruning in detail. They also demonstrate recognition accuracy versus speed trade-offs when this technique is used.

3) *Trace matching*: The final stage of recognition is matching the input data against prototypes in an ink model database. After using features for pruning the models, it is necessary to deal with only a fraction of the collection. In our system we use an elastic matching algorithm [26] to detect the closest model. We have also looked into using other trace matching algorithms. Having multiple recognizers in the system requires a reliable scheme to combine the ink recognition results. One solution is the simple voting method used in InftyEditor [7], or weighted voting as implemented in MathBrush [12].

4) *Combining recognition and prediction*: In the present version of the Mathink system, we use character prediction to assist the recognizer. We have built a mathematical context dictionary, based on observed symbol frequencies in mathematical expressions encountered in practice, as described in [24]. Character prediction is then derived from the most likely continuation of the partially entered formula. Recognition and prediction results are combined, producing a final ranking of recognition candidates. The diagram given in Figure 3 shows how recognition results and context information are used together. Another way of thinking about character prediction is to view it as a recognizer that ignores the current trace data and relies on history alone. From this point of view combining a predictor with a recognizer is no different than combining two recognizers.

C. Structure Analysis

The two-dimensional nature of mathematical notation introduces challenges that do not arise in the recognition of natural language text. For mathematical pen input, as well as recognizing each individual character, we also have to interpret its position and role in expression layout. Then, based on layout information, we build an expression tree that encodes the structure of the input formula. Our approach to structure analysis uses two steps: First, an input expression is parsed to a layout tree representing spatial relations between parts of the formula. In the second pass, the layout tree is transformed to an expression tree that encodes mathematical semantics. Our

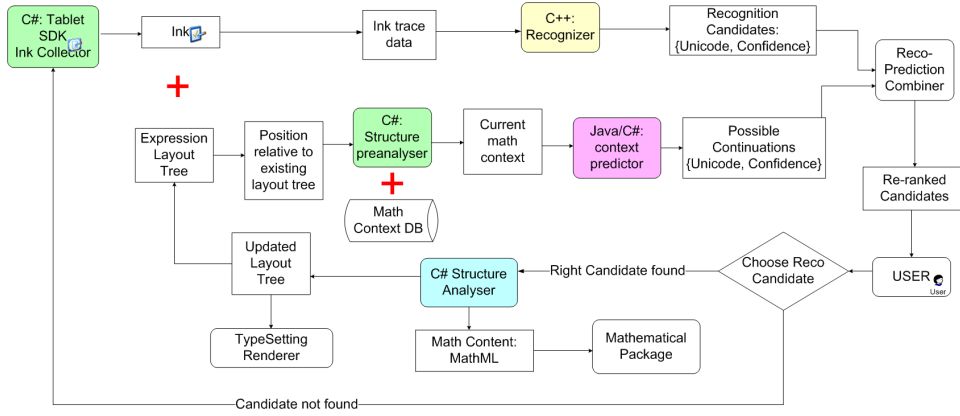


Fig. 3. Mathink recognizer work-flow using character prediction

$$\sum_{i=1}^{\infty} \frac{a^i}{i!}$$

$$a^m \rightarrow \frac{a^m}{n}$$

Fig. 5. Changes in layout organization determined by the context

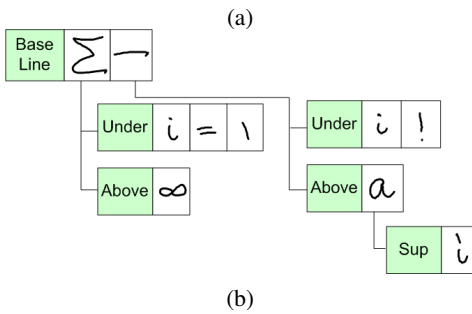


Fig. 4. An example of a layout tree

approach is related to the technique developed in [32] and used in the FFES [25] editor. However we use only two tree transformations, while [19] and [32] suggest three passes.

1) *Layout trees*: During the first pass over the handwritten input we neither try to detect its structural hierarchy nor to determine its mathematical semantics. At this stage we are concerned only with the relative positioning of the characters, both as ink and as the corresponding typeset rendering. The structure of the layout tree is organized according to baselines detected in the input expression. On each baseline we select a *leading character* that is used as a reference to adjust the rest of the symbols at the same base level. An example of a layout tree for the expression of Figure 4.a is shown in Figure 4.b.

As described in Section IV-B4, to calculate prediction values used in character recognition, the Mathink recognizer recomputes the mathematical context every time a new character is written. As can be seen from the scheme of Figure 3, the context is derived from a local layout of the expression structure. This implies that the layout tree must be calculated even for partially-entered expressions, every time a new glyph is added to the formula. Therefore, in contrast to the approach

taken in [32], during this pass we try to detect fractions and we also pre-calculate superscript and subscript relations. In some cases this preliminary layout has to be re-adjusted later, as the rest of the expression is entered. Figure 5 illustrates a common situation where a superscript transforms to numerator as a fraction bar is written.

2) *Expression trees*: Expression trees represent the hierarchical structure of the formula and encode some of its mathematical semantics. A layout tree contains most of the information for the analyzer to determine the structure of an input formula. Using several heuristics, our method groups the nodes of layout trees into tokens, then assembles tokens into subexpressions. After these procedures the tree often requires reorganization to respect matching delimiters. This transformation will adjust, for example, semantic relations between subexpressions in parentheses and scripts, e.g. converting what could be written in TeX as $\$(a+b\{) ^n\}$$ to $\$((a+b)) ^n\$$.

At this stage the structure analyzer explicitly marks decimal numbers, identifiers, operators, fractions, radicals, groups, etc. Each node of the tree is assigned a role in the expression and marked with a corresponding label. The structure recognizer also attempts to distinguish function names from variable names and implicit multiplication from function application or multi-character identifiers and their juxtapositions. This is a hard problem that in general requires deep semantics and we only partially resolve it, based on a dictionary of function names and common conventions of implicit multiplication. Figure 6 shows the transformation of a layout tree (a) to the corresponding expression tree (b) after two-stage grouping, detecting fenced subexpressions and labeling.

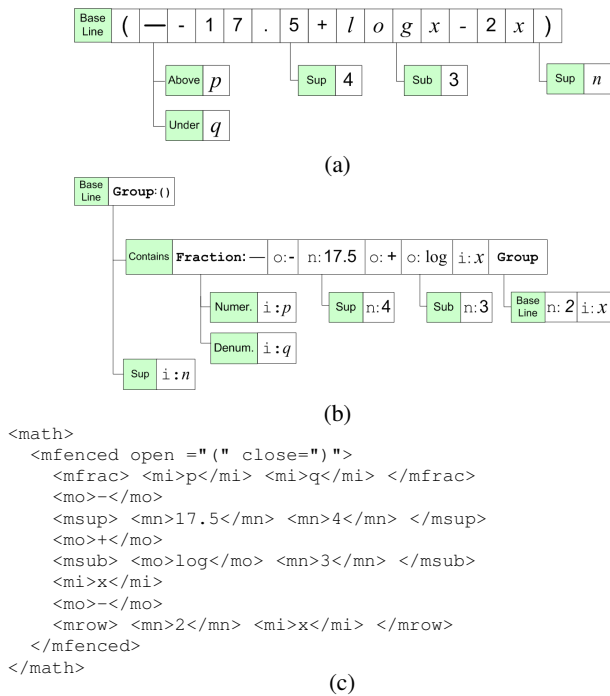


Fig. 6. Layout tree (a) and generated expression (b) and MathML (c) trees

D. Extracting Mathematical Content

The expression tree built in structure analysis already encodes some mathematical semantics. The labels on the nodes contain information about numbers, operators and variables. Parts of the expression between matching opening and closing delimiters are organized in fenced groups. Fraction parts and radicals are also explicitly marked. These preparations allow the expression tree to be directly transformed to Presentation MathML [1]. Thus the expression tree from Figure 6.b will be converted into the equivalent MathML tree shown in Figure 6.c.

We chose not to export recognition results as \TeX because it cannot easily encode all the required semantics and MathML has become sufficiently widespread. Additional advantages of MathML are platform and application independence and ease of conversion to other formats, including \TeX [20].

V. CONNECTION TO MATHEMATICAL ENGINES

After digital ink has been collected and processed, and the mathematical expression is recognized and parsed to a standard format, the pen-based application moves to the next stage, where the user may want manipulate the resulting expression. This manipulation will, in general, imply non-trivial mathematical transformation. This requires providing a connection to a mathematical engine from the handwriting recognition environment. In this section we present two approaches that allow our Mathink component to communicate with the MAPLE symbolic computation package. A similar approach would allow connection with other computer algebra systems (CAS).

A. Using the OpenMaple Interface

To send a request for computation to MAPLE from another application, we need to access the MAPLE kernel from outside the CAS shell. Each request must be expressed as a valid Maple command with all of its arguments encoded in a Maple-compatible format. For this, we use supplementary tools offered by the MAPLE package: Starting with version 9.0, the OpenMaple interface allows access to the computer algebra system functionality via API calls from C++, Java or Fortran. Furthermore, a built-in parser for both Content and Presentation MathML allows importing mathematical data encoded in these formats into MAPLE.

Once the handwritten expression is translated to MathML, it can be used in computation via OpenMaple API calls. To enable support this option in the Mathink environment we have developed a special CAS-communication module. It starts the MAPLE kernel as a background process, wraps the user's request in a MAPLE instruction and passes it along with MathML arguments using the OpenMaple protocol to the MAPLE kernel. The results of the computation are exported as Presentation MathML, sent back to the Mathink system and displayed.

This approach is similar to that used in the MathBrush [12] system, though Mathink has an independent implementation. Currently the Mathink interface allows transmission of recognition results for numeric evaluation, symbolic calculation, simplification and factorization. Nothing new would be required to support a variety of other operations. Mathink does not yet support equation solving, partial differentiation or other operations that require additional semantic analysis to detect variable instances in the the input. By identifying variable names the user could be offered a choice of "solve for .." or "differentiate with respect to .." operations.

B. A Mathink Control as an Internal Maple Component

While using the OpenMaple API is relatively easy to implement and natural to use, it has certain restrictions. These are limitations in accessing the MAPLE kernel from an outside application and ambiguities arise in conversion of mathematical notations to semantic content. Moreover, as a practical matter, the built-in MAPLE parser for Presentation MathML is not error-free.

To overcome these difficulties we introduced an alternative architectural solution that allows the pen-based component to be run inside the MAPLE environment. As discussed in Section III-B, it is possible to export the Mathink component as a Java object which can then be integrated with MAPLE GUI code. From there, the Mathink component has direct access to the MAPLE kernel. This allows the Mathink control to enjoy a larger range of computer algebra system functionality and it can form more complex and detailed requests by using native MAPLE instructions. Furthermore, hosted inside MAPLE, Mathink can operate with mathematical objects at an internal level of the symbolic engine. This includes direct translation of recognized input to MAPLE structures, as shown in Figure 7.

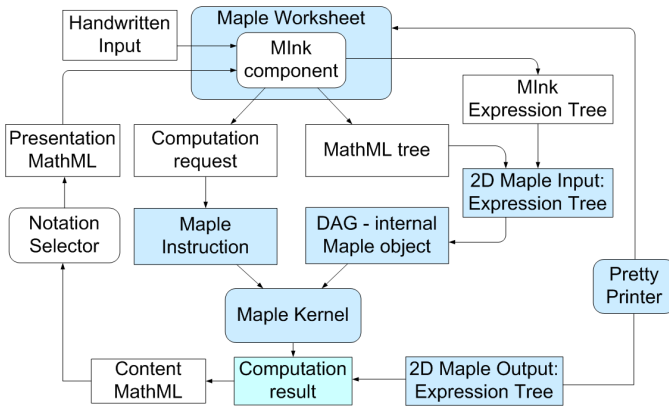
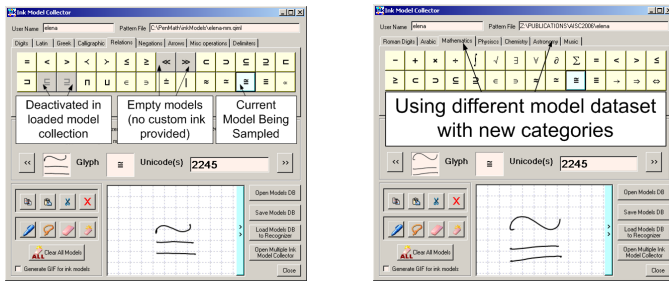


Fig. 7. Mathink accessing Maple kernel as an internal Maple component



(a) Dataset 1 (b) Dataset 2
Fig. 8. Ink Model Collector customizable by user

VI. CUSTOMIZING THE MATHINK SYSTEM

The Mathink system provides several ways to customize its behaviour. This is both for user convenience and to allow easier experimentation with user interface design. Some of the principal areas of customization are personalizing ink model collection and dynamic addition of user-created models, which we describe in this section.

A. Personalized Model Collections

1) *Customizable alphabets*: While most existing recognizer systems come with a pre-defined set of characters and glyphs that they can process, we have taken a different approach. The Mathink system allows the user to specify which mathematical symbols and glyphs he or she is going to use. This allows a wide range of characters to be used without sacrificing accuracy. The user has control over the exact number, type and appearance of glyphs known to the recognizer. The Mathink installation kit comes with a number of default model sets that are extensible by the user via adding new categories and populating them with models. Figures 8.a and 8.b show the interface for training the recognizer on different datasets.

2) *Training the recognizer on user ink models*: After the user has specified the subset of mathematical glyphs for the recognizer to deal with, the system may either be trained or default models can be used. Training consists of providing personalized ink examples for every model. Even though the default model set covers more than 200 mathematical glyphs and provides generic handwritten samples for each, the




Fig. 9. Different styles of writing the letter β

user's style of writing may differ from the default prototypes. Therefore, the system provides an interface for collecting ink samples.

a) *Multiple model sampling*: A particular symbol may often be written in a number of different ways, varying in the number of strokes, their order and direction. These criteria are used to distinguish different *writing styles* for a character. For example, one may write the Greek letter β in the three different manners shown in Figure 9. Instead of forcing the user to choose a single style for each character, we collect all of the styles the user provides during the training session. Then all these different styles are loaded into the recognizer, and all are assigned the same glyph value.

Within the same style, models may vary, having different curvatures and inconsistent cusps or intersections. Figure 10 shows the process of collecting samples of the letter β , all written in the same manner. In this case all strokes the strokes are combined together to form an *average ink model* that reflects the consistent features of the nine samples.

b) *Computing an average model*: Instances of the same model may have quite different data, for example containing a different number of points in corresponding strokes. Therefore, to compute the average model we re-sample the stroke traces. The average model calculation also tries to match vertical and horizontal cusps of all models. In case a cusp is detected in one model, but not found in another, a local extremum is used instead. The average model is updated dynamically, as the user writes. This helps to detect if one or more samples are unintentionally written in a style different from the one being collected. The user does not have to erase "trouble" samples to refine the average model: any sample can be excluded from the calculation by deactivating its frame.

c) *Ink examining utilities*: Sometimes average models may contain undesirable features, such as unexpected or misplaced cusps, loops or intersections. In this case the user may wish to have a closer look at the model to detect the source of the problem. Mathink provides a built-in *ink examining tool* to zoom in on the model and to see the ink trace computed for each average stroke. Writing direction and stroke order can be seen easily from the labels of each point in the trace. Moreover, the ink examining tool helps to prevent blind acceptance of inappropriate average models as ink prototypes, and therefore, helps to ensure the integrity of the model databases.

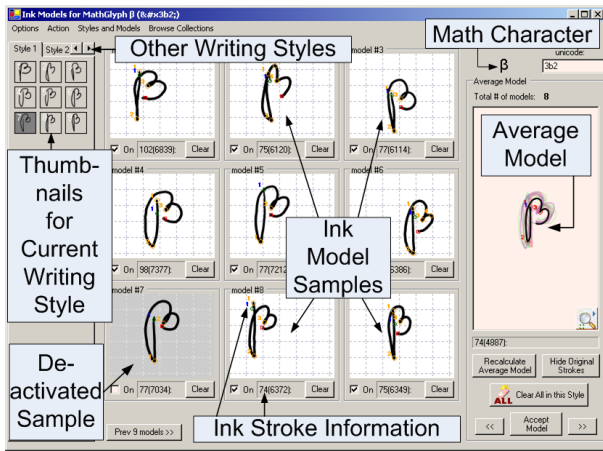


Fig. 10. An interface for collecting ink samples in different writing styles

B. Dynamic Addition of New Models

Sometimes, while working with the Mathink application, a user may need certain symbols not present in the model database. Adding a new glyph model to the dataset is a straight-forward process. However, editing a configuration file, then providing ink prototypes for new model and finally, uploading an updated model collection to the recognizer is time-consuming, and is especially inconvenient if the user needs to add several new entries at different times.

To make this process easier, the Mathink system offers the feature of *dynamic model addition* to allow quick introduction of new glyphs without re-loading main dataset. When the user writes a new symbol unknown to the recognizer, in addition to the recognition candidates an “Add new model” button is given. Clicking on this button brings up a form already containing the ink strokes that were not recognized. The user can either enter the corresponding character or characters using the keyboard or choose one of the Unicode symbols from a Unicode table, as shown in Figure 11. The same New Model form is also accessible from the main toolbar, so a user can add new models at any time. Selected strokes or the most recent unrecognized ink is automatically inserted into the ink collection box of the form. A model browser is provided, so the user can add several new entries without being forced to close and re-open the window after entering each model.

A collection of these added models can be saved and manually loaded in a subsequent session or they may be permanently added to the main model dataset. Storing newly entered models separately from the main model collection has the advantage of customizing recognizer sessions on a more fine-grained level. Loading a small patch on top of a general model collection allows, for example, to use notational preferences and avoid ambiguities when the Mathink system is used in domains with specialized notations. In this manner, the recognizer can be aware of glyph models specific to the current session, but uncommon in the user’s general writing and therefore not loaded by default.

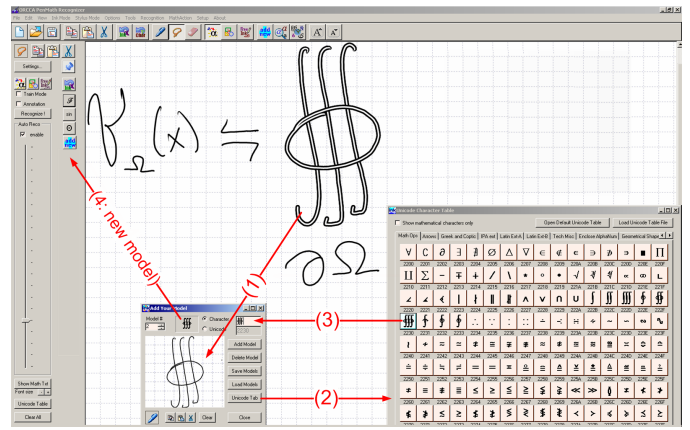


Fig. 11. Dynamically adding an ink model

VII. TESTING AND ANNOTATION TOOLS

In addition to offering a mathematical pen-based interface, the Mathink system provides testing environments to conduct experiments and adjust recognition behaviour. These include tools to collect large sets of handwritten samples, facilities for ink annotation and automated performance testing. We briefly describe each of these.

A. Collecting Test Data

In order to evaluate the performance of the mathematical recognizer, we have tested it on a number of handwritten samples provided by a set of different writers. We have collected samples of handwritten expressions for a fixed set of mathematical formulæ. To facilitate the process of obtaining handwritten samples from users, we have created several collections of test formulæ, classified by the scientific domain, formula size and structural complexity. The Mathink system, run in *sample collection mode* can load a desired set of test formulæ and show them to the writer one at a time. The user then provides a handwritten sample for the formula currently displayed. This is shown in Figure 12. When the writer finishes entering an expression and moves to the next one, the handwritten data is stored along with the reference to the original expression.

In our case we have 170 tests that we verify whenever the Mathink system is changed. Automated testing of recognition performance requires ground truth annotation to compare with recognizer output. To do this we annotate collected ink with ground truth and an interface is provided for this purpose.

B. Ink Annotation Tool

As discussed in Section IV, the recognizer uses context information to derive the total recognition confidence for each recognition candidate. To eliminate possible errors in context calculation during test runs, we also store the genuine context information for every ink glyph along with its character value.

To assign the ground truth to each ink character in a test file, we annotate ink in a semi-automated process supervised by a human user. The simplest way to provide ink annotation is to run the system recognizer on the whole expression, and then to

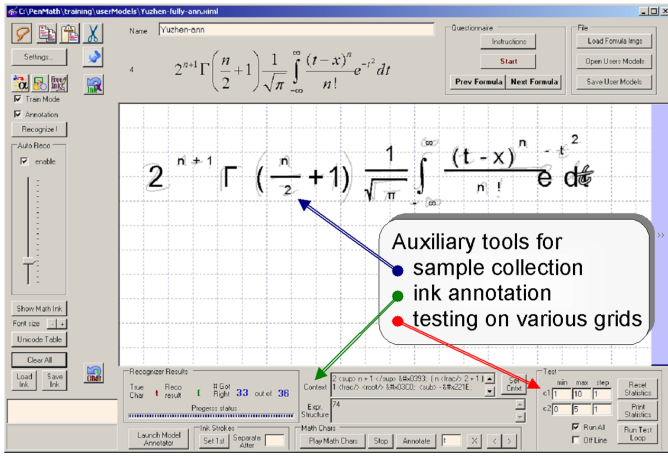


Fig. 12. Mathink application in training and testing mode

correct mis-recognized characters. Manual correction is done by clicking on an incorrectly recognized character and entering its true value from the keyboard or using the Unicode pallets. In addition, the context for any character in the expression can be adjusted using the annotation panel. This method of ink annotation is suitable only for expressions with a high rate of recognition accuracy. If the recognizer gets more than a few characters in an expression wrong, we can use the second method, which we call *annotation copying*.

The annotation copying tool transfers information from an already annotated expression to one being processed. A fully annotated expression is opened in a *master annotation window*, then the character values along with the context are copied between matching ink entries. This method is especially useful for annotating a large collection of samples provided by users for a fixed set of formulæ. In this case, for each formula in the questionnaire only one handwritten sample has to be manually annotated and the rest can be simply copied from the prototype. The reason this process cannot be done completely automatically is because there are different orders in which users write certain expression entries, such as double scripts, fractions, parentheses and radical expressions. This affects the context information: We have to take into account in which order the user wrote the expression and change the context correspondingly. For example, with the expression $\sum_{i=1}^{N^2+1} F_i(x)$, if the user writes the upper limit first, then the context for N should not contain $\sum_{i=1}^{N^2+1}$, but simply $\sum_{i=1}^{N^2+1}$.

C. Testing to Find Optimal System Settings

The character recognizer and the structure analyzer depend on a number of tunable parameters, such as threshold values used in baseline and script detection, minimum confidence for a character to be considered as a potential recognition candidate, maximum number of candidates to return from the recognizer to the structure analyzer, *etc.* For some of these parameters a default value can be relatively easily estimated and adjusted after a small number of test runs. There are, however, parameters that are impossible to approximate without having

results of practical testing on a large amount of data. Weight coefficients used in combination of the recognition confidence with prediction certainty are examples of such parameters. We have decided to search for good combinations of these coefficients experimentally, as we have described in [24].

To conduct these experiments, an automated testing facility was added to the Mathink environment, allowing the recognizer to run on handwritten samples, while different parameter values are used. A background process, collecting statistics on best combinations gives us a range of acceptable values. To give an idea of the size of this experiment for parameter estimation, we mention that each handwritten expression containing from 10 to 40 characters, was run on 10×11 parameter grid, once for each of several combining methods. Given that each writer provided 17 handwritten expression, we had to run the character recognizer about 84150 times to process all samples from one user, which took about 23 hours for each of the writers.

VIII. CONCLUSIONS

We have studied methods to organize pen interfaces for mathematical computing in a variety of environments. We have suggested an architectural approach that ensures portability of a pen-based frameworks across platforms and hosting applications. We have demonstrated how this solution can be instantiated on Windows platforms, using the Tablet PC SDK for high-quality ink processing.

We have presented a software package, Mathink, for experimentation with mathematical handwriting recognition. Mathink can be used as pen-based front-end to computer algebra systems and document editors. We have shown how our Mathink system has been used to annotate an ink collection with ground truth and how this can then be used to test and adjust the recognizer module to achieve better performance.

ACKNOWLEDGMENTS AND NOTES

We thank Xiaofang Xie for the implementation of the character recognizer for the Mathink system, Yu-Hong Wang from Maplesoft for his assistance in integrating Mathink interface into MAPLE computer algebra package and Sandra Huerter for her editorial assistance.

The work presented here was part of a larger project also involving investigators at the University of Waterloo. The Waterloo group developed a separate software testbed, MathBrush, which is described elsewhere [12].

An earlier version of this paper was presented orally at the conference *Communicating Mathematics in the Digital Era*, held 25-18 August 2006 in Aviero, Portugal, and appeared as an ORCCA technical report [22]. The present article differs by some editorial corrections, amendments for clarity and addition of new references.

To preserve the integrity of the article we have not updated the text to reflect subsequent advances in our approach. The principal difference is that our current approach provides faster character recognition, device independence and avoids resampling by using a truncated orthogonal series representation for strokes [4], [8], [9], [11].

REFERENCES

- [1] R. Ausbrooks, S. Buswell, D. Carlisle, S. Dalmas, S. Devitt, A. Diaz, M. Froumentin, R. Hunter, P. Ion, M. Kohlhase, R. Miner, N. Poppelier, B. Smith, N. Soiffer, R. Sutor, and S. Watt. *Mathematical Markup Language (MathML) Version 2.0 (Second Edition)*. World Wide Web Consortium, 2003.
- [2] R. Avitzur. Your own handprinting recognition engine. *Dr Dobbs Journal*, 17(4):32–37, 1992.
- [3] *The AWT Native Interface*. Sun Microsystems, http://java.sun.com/j2se/1.4.2/docs/guide/awt/1.3/AWT_Native_Interface.html, 1999.
- [4] B. W. Char and S. M. Watt. Representing and characterizing handwritten mathematical symbols through succinct functional approximation. In *Proc. Int'l Conf. on Document Analysis and Recognition (ICDAR 2007)*, pages 1198–1202. IEEE Computer Society, 2007.
- [5] Y.-M. Chee, M. Froumentin, and S. M. Watt, editors. *Ink Markup Language (InkML)*. World Wide Web Consortium, <http://www.w3.org/TR/2006/WD-InkML-20061023>, 2006.
- [6] *COM Interop*. MSDN Library, <http://msdn.microsoft.com/library/default.asp?url=/library/en-us/csref/html/vcwlkcominteroppart1clienttutorial.asp>, 2003.
- [7] M. Fujimoto, T. Kanahori, and M. Suzuki. Infty Editor—a mathematics typesetting tool with a handwriting interface and a graphical front-end to OpenXM servers. In *Computer Algebra—Algorithms, Implementations and Applications, RIMS*, volume 1335, pages 217–226, 2003.
- [8] O. Golubitsky and S. M. Watt. Online stroke modeling for handwriting recognition. In *Proc. 18th Annual Int'l Conf. on Computer Science and Software Engineering (CASCON 2008)*, pages 72–80. IBM Canada, 2008.
- [9] O. Golubitsky and S. M. Watt. Online computation of similarity between handwritten characters. In *Proc. Document Recognition and Retrieval XVI (DRR XVI)*. (to appear), 2009.
- [10] B. Keshari and S. M. Watt. Hybrid mathematical symbol recognition using support vector machines. In *Proc. Int'l Conf. on Document Analysis and Recognition (ICDAR 2007)*, pages 859–863. IEEE Computer Society, 2007.
- [11] B. Keshari and S. M. Watt. Online mathematical symbol recognition using svms with features from functional approximation. In *Proc. Mathematical User-Interfaces Workshop 2008 (MathUI 2008)*, 2008.
- [12] G. Labahn, S. MacLean, M. Marzouk, I. Rutherford, and D. Tausky. A preliminary report on the MathBrush pen-math system. In *Maple Conf. 2006*, pages 162–178, 2006.
- [13] *Maple User Manual*. Maplesoft, 2005.
- [14] R. Marzinkewitsch. Operating computer algebra systems by handprinted input. In S. M. Watt, editor, *International Symposium on Symbolic and Algebraic Computation*, pages 411–413. ACM Press, 1991.
- [15] E. Miller, N. Matsakis, and P. Viola. Learning from one example through shared densities on transforms. In *IEEE Conf. on Computer Vision and Pattern Recognition*, volume 1, pages 464–471, 2000.
- [16] *Platform Invoke Services*. MSDN Library, <http://msdn.microsoft.com/library/default.asp?url=/library/en-us/csref/html/vcwlkPlatformInvokeTutorial.asp>, 2003.
- [17] D. Prusa and V. Hlaváč. Mathematical formulae recognition using 2D grammars. In *Proc. Int'l Conf. on Document Analysis and Recognition (ICDAR 2007)*, pages 849–853. IEEE Computer Society, 2007.
- [18] T. H. Rhee, K.-E. Kim, and J. H. Kim. Robust recognition of handwritten mathematical expressions using search-based structure analysis. In *Proc. Int'l Conf. on Frontiers in Handwriting Recognition (ICFHR 2008)*. CENPARMI Concordia U., 2008.
- [19] I. Rutherford. Structural analysis for pen-based math input. Master's thesis, University of Waterloo, 2005.
- [20] E. Smirnova and S. M. Watt. *MathML to LaTeX on-line converter*. <http://www.orcca.on.ca/MathML/texmml/mmltotex.html>, 2002-2005.
- [21] E. Smirnova and S. M. Watt. A context for pen-based computing. In *Maple Conf. 2005*, pages 409–422. Maplesoft, 2005.
- [22] E. Smirnova and S. M. Watt. A pen-based mathematical environment: Mathink. Technical Report TR-06-05, Ontario Research Centre for Computer Algebra, 2006.
- [23] E. Smirnova and S. M. Watt. Aspects of mathematical expression analysis in arabic handwriting. In *Proc. Int'l Conf. on Document Analysis and Recognition (ICDAR 2007)*, pages 1183–1187. IEEE Computer Society, 2007.
- [24] E. Smirnova and S. M. Watt. Context-sensitive mathematical character recognition. In *Proc. Int'l Conf. on Frontiers in Handwriting Recognition (ICFHR 2008)*, pages 604–610. CENPARMI Concordia U., 2008.
- [25] S. Smithies, K. Novins, and J. Arvo. A handwriting-based equation editor. In *Int'l Conf. on MathML and Math on the Web, (MathML 2002)*, 1999.
- [26] C. Tappert. Cursive script recognition by elastic matching. *IBM Journal of Research Development*, 26(6):765–771, 1982.
- [27] *The Unicode Standard, Version 5.0*. Addison-Wesley, 2007.
- [28] B. Wan and S. M. Watt. An interactive mathematical handwriting recognizer for the Pocket PC. In *Int'l Conf. on MathML and Math on the Web, (MathML 2002)*, 2002.
- [29] S. M. Watt. An empirical measure on the set of symbols occurring in engineering mathematics texts. In *Proc. 8th IAPR International Workshop on Document Analysis Systems (DAS 2008)*, pages 557–564. IEEE Computer Society, 2008.
- [30] S. M. Watt and X. Xie. Prototype pruning by feature extraction for handwritten mathematical symbol recognition. In *Maple Conf. 2005*, pages 423–437. Maplesoft, 2005.
- [31] S. M. Watt and X. Xie. Recognition for large sets of handwritten mathematical symbols. In *IEEE Int'l Conf. on Document Analysis and Recognition (ICDAR 2005)*, pages 740–744, 2005.
- [32] R. Zanibbi, D. Blostein, and J. Cordy. Recognizing mathematical expressions using tree transformation. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 24, pages 1455–1467, 2002.
- [33] R. Zanibbi, K. Novins, J. Arvo, and K. Zanibbi. Aiding manipulation of handwritten mathematical expressions through style-preserving morphs. In *Graphics Interface*, pages 127–134, 2001.
- [34] L. Zhang and R. Fateman. Survey of user input models for mathematical recognition: Keyboards, mice, tablets, voice. Technical report, University of California, 2003.