

作りましょう 0.11
パラメタ方式フォントファミリ
ユーザマニュアル

Tsukurimashou 0.11
Parametric Font Family
User Manual

Matthew Skala
mskala@ansuz.sooke.bc.ca
2021年7月1日 July 1, 2021

This project's English-language home page is at

<https://tsukurimashou.osdn.jp/index.php.en>.

このプロジェクトは、日本語のページが

<https://tsukurimashou.osdn.jp/index.php.ja>です。

User manual for Tsukurimashou

Copyright © 2011, 2012, 2013, 2014, 2015, 2017, 2021 Matthew Skala

This program is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, version 3.

As a special exception, if you create a document which uses this font, and embed this font or unaltered portions of this font into the document, this font does not by itself cause the resulting document to be covered by the GNU General Public License. This exception does not however invalidate any other reasons why the document might be covered by the GNU General Public License. If you modify this font, you may extend this exception to your version of the font, but you are not obligated to do so. If you do not wish to do so, delete this exception statement from your version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program. If not, see <http://www.gnu.org/licenses/>.

目次 Contents

イントロ Introduction	5
0.11のニュース What's new in 0.11?	10
外のプロジェクト Other similar projects	11
デベロップメントロードマップ Development roadmap	15
外のソフトのバグ Relevant bugs in other software	17
METAFONTなど METAFONT, etc.	17
フォントフォージェ FontForge	19
X _Y TEXX _Y TEX	22
外のソフトウェア Other software	24
『作りましょう』の使い方 Using Tsukurimashou	26
OpenTypeのフィーチャー OpenType Features	27
外に分すう Alternate Fractions (afrc)	28
キャプがスモールになって Capitals to Small Caps (c2sc)	29
文脈代替字 Contextual Alternates (calt)	29
グリフの併合と分解 Glyph (De)Composition (ccmp)	29
合字 Ligatures (liga)	30
文字に付け方 Mark to Base Positioning (mark)	31
点に付け方 Mark to Mark Positioning (mkmk)	31
メタデータ Metadata Table (name)	32
巴の花形 Ornaments (ornm)	32
FontAnvilだけのメタデータ FontAnvil-specific Metadata (pfed)	32
スモールキャピタル Small Caps (smcp)	34
ヘビーメタルウムラウト Heavy Metal Umlaut (ss01)	34
丸つき字 Enclosed Letters and Numerals (ss02)	35
私用符号位置を避けて Limit PUA Code Points (xpua)	36
付加文字 Extra glyph coverage	38
げんじもん Genjimon	38
えききょう I Ching	38
解字の文字 Ideographic Description Characters	39
公のユーログリフ Official Euro Sign	39
ハングル語 Korean Language Support	41

フォントの名前について Regarding Font Names	42
『作りましょう』を作りましょう! Building Tsukurimashou	47
ビルドのシステム Build system	49
寄生パッケージ Parasite packages	55
バージョン管理システムからビルド Building from a version-control checkout	58
頒布のパッケージを作って Building distribution packages	61
ツールの文書化 Tool documentation	64
ビルドのシステムの図示 Build system diagrams	68
ハムログ Hamlog	69
カーニングしかた Kerning	78

イントロ

Introduction

日本ごのユーザマニュアルをまちうけば、ごめんなさい。みらいは、日本ごのユーザマニュアルをかきます。

I want to learn Japanese. That's a large project, likely to involve years of daily study to memorize a few tens of thousands of words. It's clearly within the range of human instrumentality, because many millions of Japanese people have done it; but most of them started in infancy, and it's widely believed that people's brains change during childhood in such a way that it's much harder to learn languages if you start as an adult. I also would like to be fully literate within somewhat less than the 15 years or so that it takes a native learner to gain that skill.

What about designing a Japanese-language typeface family? Typeface design is a difficult activity, requiring hours of work by an expert to design each glyph. A typeface for English requires maybe 200 glyphs or so; one for Japanese requires thousands. That makes the English-language typeface about a year's full-time work, and the Japanese-language typeface enough work (20 to 25 man-years) that in practice it's rare for such a project to be completed by just one person working alone, at all. There is also the minor detail that I said a "typeface family"—one of those typically consists of five or six individual faces, so the time estimate increases to between 100 and 150 years.

However, I've already decided to spend the time on the language-learning project. And I'm sure that if I design a glyph for a character, I'm going to remember that character a lot better than if I just see it a few times on a flash card. And as a highly proficient user of computer automation, I have access to a number of efficiency-increasing techniques that most font designers don't. So the deal is, I think if I study the language **and also design a typeface family for it**, I might be able to finish both big projects with less, or not much more, effort than just completing the one big project of learning the language alone. And then I'd also end up with a custom-made Japanese-

language typeface family, which would be a neat thing to have. And so, here I am, building one.

Because this is a parametrized design built with Metafont, users can generate a potentially unlimited number of typeface designs from the source code; but I've provided parameters for several ready-made faces, each of which can be drawn in monospace or proportional forms. The mainline Tsukurimashou fonts are intended for Japanese, but version 0.5 also introduced a series of fonts called Mandeubsida, designed for Korean hangul. There is also some experimental code intended to eventually become a set of blackletter fonts, but it is disabled by default and not currently usable.

The main goal of this project is for me to learn the kanji by designing glyphs for them, and so far it does seem to be helping my learning process, though the nature of computer programming is certainly having some unusual disruptive effects. For instance, in some cases I appear to have unintentionally memorized the Unicode code points of kanji without learning their meanings or pronunciations. I may end up learning to speak Japanese just like a robot; but by the time I'm fluent, the Japanese demographic collapse will be in full swing and they'll have replaced much of their population with humanoid interfaces anyway, so maybe I'll fit right in.

Please understand that the finished product is not the point so much as the process of creation, hence the name. Furthermore, although the fonts cover the MES-1 subset of Unicode, and thus can in principle be used for almost all popular languages that use the Latin script, and I know that at this stage most users are probably more likely to use the fonts for English than for anything else, nonetheless these fonts are intended for eventual primary use with the Japanese language. Some decisions on the Latin characters were driven by that consideration—in particular, the simplistic serif design and weighting in the Latin characters of Tsukurimashou Mincho, and the limited customization of things like diacritical mark positioning in Latin characters not used by English. I biased some marks (notably ogonek and haček) toward the styles appropriate for Czech and Polish as a nod to my own ancestors, but I cannot read those languages myself, and I cannot claim that these

fonts will really look right for them. I'm not interested in spending a lot of time tweaking the Latin to be perfect because it's not really the point. I already know how to read and write English.

Some other notes:

- ☯ Since my learning the kanji is a big part of the goal of this project, “labour-saving” approaches that would relieve me of having to look at all the kanji individually myself (for instance, by feeding a pre-existing database of kanji shapes into my existing general font technology) are not appropriate to the original intention of the project.
- ☯ As of version 0.9 (August 2014) I'm not happy with the weight of punctuation marks, nor the sizes of serifs in the Mincho Latin character range. Both these things will probably change in future versions, with consequent changes in spacing.
- ☯ I would like to include at least some support for vertical script, but it is not a high priority. One obstacle is that I don't have access to competent vertical typesetting software, whether the font could support it or not.
- ☯ Tsukurimashou is designed primarily for typesetting Japanese, secondarily for English. I have no immediate plans to support other Han-script languages (such as any dialect of Chinese) nor put a lot of effort into tweaking the fine details of characters only intended for use in occasional foreign words.
- ☯ Support of Korean is limited because of my limited knowledge of that language; and learning it is not a high priority for me. At this point the Mandeubsida fonts only support Korean hangul, not hanja (which are the Korean equivalent of kanji, but just different enough that copying over the Tsukurimashou kanji would not be good enough).
- ☯ I reserve the right to add features that I think are fun, even if they are not useful.

- ☹ Tsukurimashou is designed as a vector font, assuming an output device with sufficient resolution to reproduce it. In practice, that probably means a high-quality laser printer. I have not spent time optimizing it for screens or low-resolution printers, and the hinting is automated.
- ☹ If it turns out to be too much work after all, I might abandon the whole project.
- ☹ Both building and use of Tsukurimashou require working around many bugs in third-party packages, some of which were mentioned here in earlier versions of this document. The list has now grown so long it needs its own subsection, which starts on page 17 of this document.

The Tsukurimashou fonts are distributed under the GNU General Public License, version 3, with an added paragraph clarifying that they may be embedded in documents. See the files COPYING and COPYING.GPL3, and note the following addition:

As a special exception, if you create a document which uses this font, and embed this font or unaltered portions of this font into the document, this font does not by itself cause the resulting document to be covered by the GNU General Public License. This exception does not however invalidate any other reasons why the document might be covered by the GNU General Public License. If you modify this font, you may extend this exception to your version of the font, but you are not obligated to do so. If you do not wish to do so, delete this exception statement from your version.

The license means (and this paragraph is a general summary, not overriding the binding terms of the license) that you may use the fonts at no charge; you may modify them; you may distribute them with or without modifications; but if you distribute them in binary form, you

must make the source code available. Furthermore (this is where font-embedding becomes relevant) embedding the font, for instance in a PDF file, does not in itself trigger the source-distribution requirement.

My plan is that at some point in the future, when the fonts are in a more useful and complete form, I will make precompiled binaries available through commercial online channels. That will serve several purposes: it will allow me to make some money from my work, and it will also probably encourage some people to use the fonts who wouldn't otherwise. One of the bizarre aspects of human behaviour is that some people will buy a product they would not accept for free. Okay, whatever; in such a case I'm happy to take the money for it. Having a pay option will also give anybody who wants to support my efforts, an easy way to do that. For now, though, I am distributing Tsukurimashou only as this source package, with precompiled versions included for the Kaku and Mincho styles. The Mandeubside fonts, which don't need the in-progress kanji glyphs, may also be available in precompiled form as a separate package. If you want other styles, you'll have to compile them yourself or get them from someone who has done so. This limitation is deliberate: with the fonts in their current partial form, I'd rather limit their circulation to hobbyists.

I maintain several other free font projects, which as of version 0.7 have largely been merged into the Tsukurimashou packaging and build system because they share code. These "parasite" packages appear as subdirectories of the Tsukurimashou source distribution, and may also be available as separate packages. See the discussion in Subsection for more information about parasite packages.

This documentation file gives some notes on the build system and on how to use the OpenType features built into the fonts. Other documentation files included in the package demonstrate what the fonts look like and list the current kanji coverage. Better documentation (and some day, Japanese-language documentation) will probably appear in a later version; at the moment, I'm just more interested in designing fonts than in writing about them. Of course, all the typesetting in this manual is done with fonts from this package.

The name “Tsukurimashou” could be translated as “Let’s make something!”

From time to time, people ask how they can help with the project. The one thing this project most needs is your **attention**. I am not looking for code contributions because of the pedagogical goal (part of the point is for me to learn kanji by doing all the work); I am not looking for money unless it’s enough money to reduce my need for other paid employment; but public and third-party attention would be very helpful to the project’s long-term success. Use it, share it, and above all write about it. At this point I’m not even sure anybody but me has ever so much as run the Tsukurimashou build process, and that really hurt my motivation to continue maintaining it as a published free software package when other demands in my life pile up.

The home pages for this project, where you can download the latest releases, browse the source-control repository, and so on, are:

<https://tsukurimashou.osdn.jp/index.php.en> (English)

<https://tsukurimashou.osdn.jp/index.php.ja> (日本語)

よろしくおねがいします。

Matthew Skala

mskala@ansuz.sooke.bc.ca

July 1, 2021

0.11のニュース What’s new in 0.11?

It’s been about three and a half years since the version 0.10 release, and the COVID-19 pandemic, still ongoing as I write this in June of 2021, affected more than a year of that time. I am still running my synthesizer business, North Coast Synthesis Ltd. Please visit <https://northcoastsynthesis.com/> for high-quality Eurorack modules and kits. The business is bringing in enough money for me to live on, but just barely. I’m still living from one customer order to the next.

My roadmap for Tsukurimashou is still much as it was in the last release. This version 0.11 covers all the grade-school kanji up to a little

past the halfway mark of Grade Six, now including the 20 “prefectural name” characters that the Japanese school system added to Grade Four after the time when I set up my checklists. I plan to complete Grade Six in version 0.12. It only calls for 81 more kanji, plus whatever other ones are convenient to add along the way. Then, after that, the next release will be version 1.0, with no specific additional kanji planned, but the chance to tidy up any issues that might be found in version 0.12 and have a clean release that is complete in some important sense.

I feel pretty well committed to seeing this project through to the 1.0 release, but I don't know if I will continue it past there. It has been about ten years' work now, and it will be at least a little more than that by the time 1.0 is actually released. A lot has changed in my life over that time. I haven't been as successful at learning the Japanese language as I'd hoped, and some of the incentives for me to keep trying no longer exist. This project also hasn't attracted as much third-party attention as I had hoped it would, and third-party attention is now even more important as a reason to continue it, than was earlier the case.

But having a font family that can reasonably be used for at least some practical typesetting, not just demonstration purposes, will be an important milestone and will open up a lot of possibilities. I still have some hope that once I push the version 1.0 release I can expect it to be actually adopted by people other than myself, and that will improve the profile of the project. I'm also still looking out for ways I might be able to gain money or professional status from this project, which are the things that I'd really need to have to make it feasible to keep working past 1.0.

よろしく願います。

外のプロジェクト Other similar projects

Maybe you shouldn't use this package! It is designed for specific purposes that are relevant to its designer, and although I certainly hope others will find it useful, my goals may or may not be in line with

yours. Also, although I sometimes describe Tsukurimashou as the first parameterized METAFONT family with Japanese-language coverage, that claim requires careful qualification because many projects with similar aims have existed in some form for a long time. Here are some others, going back a few decades, that you might want to check out.

This is not intended to be a complete list; in particular, I'm leaving out many sources of CJK fonts that are not METAFONT-related, and many academic papers that are not associated with publically-available fonts. There is also no doubt a great deal of research and development locked up inside commercial organizations, or published in the Chinese language and thus inaccessible to me.

- ④ “LCCD, A Language for Chinese Character Design,” Tung Yun Mei, Stanford technical report STAN-CS-80-824. He built a language similar in nature to METAFONT (collaborating with Knuth and sharing ideas with the earliest versions of METAFONT) and constructed 112 kanji for use in TAOCP. It's interesting how little has changed since this early work.
- ④ “A Chinese Meta-Font,” John Hobby and Gu Guoan, paper in TUGboat 5-2, 1984. Proof of concept and discussion of some of the graphic design issues for parameterized CJK fonts. They built 140 radicals and 128 characters, using infrastructure very similar to Knuth's techniques for Latin fonts, and the traditional Chinese stroke-based analysis of characters. High-quality parameterized designs. No apparent plan to actually turn these into a usable full-coverage family; it seems to have been meant as research into the techniques only. <http://www.tug.org/TUGboat/tb05-2/tb10hobby.pdf> John Hobby has posted the source code described in the paper on his Web site at <http://ect.bell-labs.com/who/hobby/hobbygu.tar.gz>, but as he says in the enclosing Web page, the files are written in the now-obsolete METAFONT79 language and “they are of limited use because they are not compatible with today's METAFONT.”
- ④ p_TE_X, ASCII Corporation, 1987 onward: not a font project, but a T_EX engine modified to handle 16-bit character codes and using existing fonts from other systems. Very popular in Japan; to

some extent it still is, though other projects of similar nature (mostly not listed here) have gained a lot of market share in recent years.

- ④ The Quixote Oriental Fonts Project, spearheaded by Dan Hosek, announced in a paper at the TUG 1989 Conference. Intended to be a parameterized METAFONT-native family for Chinese, Japanese, and Korean. Hosek apparently had some source code in hardcopy form displayed at the conference, but I've not been able to find the code nor any subsequent discussion of the project. <http://www.tug.org/TUGboat/tb10-4/tb26hosek.pdf>
- ④ Poor Man's Chinese and Poor Man's Japanese, 1990, Tom Ridgeway: technology for displaying 24×24 bitmap fonts through METAFONT. This was not curve tracing, nor smooth scaling, but a way to actually display the dot matrix, jaggies and all. Still available in CTAN package "poorman," but considered obsolete.
- ④ JemTEX, 1991, François Jalbert: included a program called `jis2mf` which would auto-trace 24×24 bitmaps to produce non-parameterized METAFONT code. Many sources from the 1990s (for instance, a regular Usenet posting aiming to list all then-available .mf-format fonts) describe the availability of "Metafont for 61 Japanese fonts," which is the output from this program.
- ④ The CJK package, described in TUGboat at least as early as 1997, still available though no longer popular, Werner Lemberg. Not a font project but a system for typesetting CJK text in L^AT_EX under the standard 8-bit engines, getting around the encoding issues by splitting each font into many smaller virtual fonts (similar to Tsukurimashou's "page" system). Fonts for this, at least at the outset, were usually bitmap fonts imported from other systems (one popular one was 48×48); later, as free vector fonts became available, those started to be used, some of them via auto-conversion from formats like TrueType to non-parameterized METAFONT.
- ④ HanGlyph, 1997 and 2003: a language for describing Chinese characters, and support for rendering them in MetaPost and L^AT_EX. This

is intended to address the Chinese equivalent of the “gaiji” problem: how to typeset rare characters that are not included in standard fonts or encodings. The user can describe the missing character and a small font containing just that character will be automatically created and used. In principle, HanGlyph’s technology could be used to create a full-coverage font, but as of 2012 it doesn’t appear anyone has done that. Availability and licensing terms are unclear; no released code or fonts seem to be available, but there have been papers published about it. <http://www.hanglyph.com/>

- ④ IPA Mincho and IPA Gothic fonts, 2003 onward. Note “IPA” in this case stands for “Information-technology Promotion Agency,” not the “International Phonetic Alphabet,” and these fonts do not cover that IPA. Free high-quality fonts for Japanese, TrueType format, not parameterized. <http://ossipedia.ipa.go.jp/ipafont/index.html>
- ④ X_YTEX, SIL International, 2004 onward: T_EX engine extended to handle Unicode and modern font technologies. Used to compile this document, and one of the main compatibility targets for Tsukurimashou. <http://scripts.sil.org/xetex>
- ④ Hóng Zì project by Javier R. Laguna. Aimed to be a parameterized METAFONT family for Chinese. The last release, which was in 2006, contained 125 characters. No infrastructure for addressing issues like METAFONT’s 256-glyph limit, or radicals changing shape depending on their context. Probably abandoned. However, it did make several releases of code that you can still download and compile. <http://hongzi.osdn.net/>
- ④ The KanjiVG Project, coordinated by Ulrich Apel, current in 2011. Still under development, but already has basically complete coverage, and is deployed in several important applications. This is not a font family, but a database of kanji (primarily from a Japanese point of view) broken down into strokes and radicals, with some curve points and a lot of abstract information about how the strokes correspond to the traditional radical classification (so that you can automatically recognize, for instance, what 示 and 礻 have in common), stroke order, and so on. This

is a valuable resource for dictionaries and handwriting recognition systems. Some kind of supervised semi-automatic processing could probably turn it into a font, but keeping the style consistent (because the database has multiple sources), and adding the serifs and other visual information needed for styles like Mincho, would require some significant work. <http://kanjivg.tagaini.net/> See it in action in Ben Bullock's handwritten kanji recognizer at <http://kanji.sljfaq.org/draw.html>.

- ④ Character Description Language, from Wenlin Institute, Inc. Current in 2011. This is a commercial product. It is apparently (though the Web page could be clearer on the exact nature of what they're selling) a database of character descriptions similar to KanjiVG though with wider coverage, plus a binary-only rendering library, the combination available for license at some unspecified price. It says it's capable of generating MetaPost as one of its several output formats. Not clear to what extent there is parameterization, but presumably that would be in the converter rather than in the resulting MetaPost. <http://www.wenlin.com/cdl/>
- ④ Type Project Adjustable Fonts, announced in 2012. Commercial type foundry offers to provide Japanese-language corporate fonts that are "adjustable" for weight and width. It is not clear that they actually have a full set of kanji; from the description on the Web site it appears that they have a ready-made set of parameterized kana and then they will create whatever specific kanji the client wants to pay for. The parameterization technology is evidently smarter than the purest form of linear interpolation (because they have a JavaScript demo highlighting the difference) but it still seems to be within the range of what could be accomplished with, for instance, Adobe Multiple Master. <http://typeproject.com/projects/adjustable>

デベロップメントロードマップ

Development roadmap

This version contains all the kyouiku kanji (the ones taught in Japanese elementary school) through Grade 5, and 100 of the 181 in Grade 6.

The current plan is to release a minor version after each grade level of kyouiku kanji and one halfway through each grade level, which will take us up to version 0.12 at the end of Grade 6. Version 1.0 will probably be a separate version released shortly after 0.12, with a general clean-up and renovation.

There are 1026 kyouiku kanji, though the fonts already contain far more than that number of kanji glyphs because my general practice is to add other glyphs that are convenient to add whenever they come up, regardless of their level. For instance, when I added the “gate” radical 門 it was easy to add many other kanji that consist of that wrapped around an easy pre-existing kanji, even though some of those are not in common use and one, 閨, isn’t even a real kanji at all, having been created by an error in the standards process. But having just over a thousand in the main-line roadmap makes the end of the kyouiku kanji a good milestone for the first major version number.

Although this is subject to change and cancellation, my current thought is that the next major versions would be 2.0 with the jouyou kanji (taught in high school, a total of 1110 additional glyphs), and 3.0 with the jinmeiyou kanji (the “name-only” kanji, 983 additional glyphs). I don’t know how I’d break those up into minor versions, but presumably I’d aim for a similar spacing of about 100 new characters per release. At the 3.0 point, with a little over 3000 specifically-required kanji, the fonts should be basically complete in the sense of being usable to write the full Japanese language as most reasonably competent native readers know it. The project already contains a little over 2900 kanji because I’ve opportunistically added lower-priority ones that were easy to add along the way (for instance, some that are obscure in themselves but used as components of higher-priority characters). Whether I will actually continue to work on the project past version 1.0 is undecided at the moment. It will depend on how much attention the versions up to 1.0 receive.

外のソフトのバグ Relevant bugs in other software

Building Tsukurimashou (and, especially, its documentation) requires the use of some fairly advanced features of third-party software. Some of those features are not often used; as a result, I've become an unintentional beta-tester, and in some cases a maintainer, for the third-party packages. I've previously noted bugs when they come up, in code comments and the relevant parts of this document, but as of the current version, such bugs have become so numerous that it makes sense to also have a central list.

METAFONTなど METAFONT, etc.

- ☹ The “independent variable instance serial numbers” issue in Meta-post seems to no longer be significant in reasonably recent versions. As of Tsukurimashou version 0.11, I'm no longer tracking it.
- ☹ METATYPE1 sometimes runs glyph names through the METAFONT tokenizer. At this point, I don't know how essential that is to the operation of METATYPE1 or whether it can be changed. It has subtle effects that can cause problems. One issue shows up in glyph names that contain a decimal digit followed by a dot, as in “uni1100.bug”; then what gets written into the Postscript output is “uni1100bug” because that is equivalent but more canonical in METAFONT syntax. A more serious issue shows up with the glyph name “uni1100.l1” from Tsukurimashou 0.5; in the new METATYPE1 version 0.55, that gets tokenized in a context where the token “l” is a “spark,” and so the whole compilation fails. The workaround for that was to change “l” to “lj” (for “lead jamo”). A further workaround, introduced in Tsukurimashou 0.6, was to modify the bundled METATYPE1 code to save the glyph name as a string and put that in the Postscript file instead of the tokenized version. It still also parses the names as tokens, however, so there is also code to insert a bunch of underscores in the internally-used name before doing that, to reduce the chance of name collisions from this parsing.

- ☹ METATYPE1 `pen_stroke_edge` macro: as of METATYPE1 version 0.44, for reasons unknown to me if left to its own devices it will sometimes attempt to evaluate the “turning angle” of a zero vector, and then blow up. This seems to happen most often when stroking a vector in a direction of approximately 290 degrees. As of Tsukurimashou 0.6, we are bundling a macro derived from METATYPE1 version 0.55, which seems to have fixed this bug; the workaround in previous versions of Tsukurimashou has been removed.

- ☹ Not really a bug, because they warn about it in the documentation and it is a reasonable consequence of design decisions made for good reasons, but: in some cases the METATYPE1 `pen_stroke_edge` macro’s output does not well approximate the theoretical ideal shape that would be obtained by stroking the specified pen along the specified path. A perfect result is not possible because the theoretical perfect curve is not, in general, a cubic spline curve, and there are tricky topological considerations in play too. Some approximation is necessary, and the one chosen by METATYPE1 basically uses one control point in the output envelope for each control point in the input path. That may or may not be enough control points to produce a visually nice result. It seems to especially often be a problem when there is a tight curve and not actually a sharp corner in the input path, so curvature is great but not infinite, or when there is an inflection point in the input path. These cases are best simply avoided, but Tsukurimashou’s bundled version of METATYPE1 now contains added code that attempts to detect tight curves and inflection points and add extra control points automatically to reduce the likelihood of visual problems. The detection rule the code uses was chosen by trial and error educated by dimensional analysis, not by solid theoretical results, and it will probably be incorrect in some cases. The extra points and resulting topological weirdness necessitate FontForge postprocessing, and may tend to trigger bugs in FontForge, but FontForge postprocessing was already a dangerous necessity in this project for other reasons anyway. Thanks to Shriramana Sharma and the Metapost mailing list for discussion encouraging me to work on this issue.

- ☹ METATYPE1 infrastructure in general: sometimes generates paths that some software layer (possibly METATYPE1's own code) flags as “degenerate,” triggering a fatal error. Workaround is to filter things, before rendering, through the regenerate macro in intro.mp, which removes any very short path segments. Possibly related: the Fill macro will sometimes abort in response to some conditions on “turning number” that do not appear to actually be harmful. Workaround is to use our dangerousFill in intro.mp instead, which is just a copy of Fill with the error checking removed. Since the switch to bundled code based on METATYPE1 version 0.55, this seems to be a less significant issue.

フォントフォージェ FontForge

- ☹ FontForge's list of external dependencies continues to grow, as does its emphasis on Python. There has been talk of removing “native” scripting from FontForge altogether, which would be a disaster for Tsukurimashou. Rather than waiting for FontForge to become unusable and then panicking, I've forked the project. FontForge's source repository now includes a directory for “FontAnvil,” a standalone interpreter of FontForge's native scripting language. Tsukurimashou's build system as of version 0.9 will look for FontAnvil and use it in preference to FontForge if found. My hope is that this will be beneficial to both Tsukurimashou and FontForge, by providing a stable platform for Tsukurimashou's scripts and also freeing the FontForge team from the ongoing need to support features that do not fit into their marketing strategy.
- ☹ FontForge spline geometry operations, such as overlap removal and simplification: these have historically tended to be very numerically unstable, and subject to some combination of infinite loops, segmentation faults, unexplained floating-point exceptions, bizarre error messages, incorrect output, and so on. The most recent development versions seem to be relatively good; best advice is to switch to one of those, turn debugging on, and hope.

It is also sometimes possible to get past individual problems by switching between the different floating-point formats offered by FontForge's compile-time configuration: float, double, and long double. But which of those is the best choice isn't always predictable, and as of October 2012 some FontForge developers are attempting to remove the option to choose between them (which should improve stability in the long term by improving maintainability, but may harm stability in the short term). Some numerical tweaks in actual glyph outlines also exist to try to work around these issues.

- ☹ FontForge usually fails to automatically insert appropriate subtable breaks when it reads a feature file describing a large OpenType table. Workaround: our Perl code inserts explicit breaks; trial and error needed to figure out how frequent they ought to be.
- ☹ When FontForge saves an Adobe feature file, what it writes may bear only a passing resemblance to what was actually in the font, and in particular, cannot subsequently be loaded by FontForge (sometimes cannot be loaded at all, other times can be loaded but the loaded tables behave differently from the ones that were saved). Workaround: do not save feature files from FontForge. As of October 2012 there has been some interest shown by FontForge developers in addressing this issue, but it will likely be a gradual process instead of a single complete fix.
- ☹ FontForge can segfault while writing a feature file because of dereferencing a pointer first and checking whether it was null afterward. Fixed by the FontForge developers in December 2011.
- ☹ FontForge does not apply features of the "DFLT" language system as a default to Unicode ranges that have no explicit mention in the font file (Adobe spec says it should). Workaround: explicitly list language systems to cover every Unicode range we care about.
- ☹ FontForge may write up to 23 horizontal stem hints per glyph when writing a PostScript-flavoured OTF font, resulting in a PostScript stack overflow on loading if the glyph also has a non-default advance width. Went undetected for a long time

apparently because only CJK fonts are likely to have so many horizontal stems, and CJK fonts are likely to be monospace and thus won't also have per-glyph non-default advance widths. Fixed by changing the limit to 22, with a patch I wrote that the FontForge developers accepted in December 2011.

- ☹ FontForge save and then load of an SFD file has the effect of renaming the “clig” feature to “rtla” on only one of the three machines where I've tried it. May be related to the x86-64 architecture. Not reliably reproducible. Reported on mailing list November 2011, mentioned in a developer's commit message December 2011, not observed recently and probably fixed by now.
- ☹ FontForge rasterization to BDF via FreeType as opposed to whatever other code FontForge would use: sometimes produces corrupt results, which has complicated indirect consequences because the Tsukurimashou infrastructures uses BDF files as input to the auto-kerning program. As a result, when this bug bites the horizontal metrics go screwy on some fonts—notably, Mandeub-sida Batang PS ends up looking monospaced. Reported to FontForge mailing list in November 2011, inspiring FreeType maintainer Werner Lemberg to find and fix an unrelated bug in FreeType (it was unable to process BDF files with high code points; see <http://savannah.nongnu.org/bugs/?34896>). But that bug did not actually affect us because FontForge uses its own code for BDFs instead of FreeType's anyway, and Lemberg says can't help with the corrupt rasterization. Workaround is to compile FontForge without FreeType support; from the GUI it is possible to just turn off FreeType rasterization on individual bitmap-creation operations, but that option doesn't seem to be available from the scripting language and so the package has to actually be built without FreeType. As of April 2012, seems not to be an issue with Arch Linux's packaged versions. However, as of a few minutes before the Tsukurimashou 0.8 release (ouch!), this is a problem again with the latest FontForge, and the latest FontForge now requires FreeType, so the workaround of disabling FreeType no longer works. I've reported it as FontForge GitHub issue number 685 (<https://github.com/fontforge/fontforge/issues/685>), am hack-

ing my local copy of FontForge to generate distribution fonts, and will hope to get it resolved properly in mainline FontForge as soon as possible.

- ☹ FontForge attempts to modify the names of all glyphs that it thinks are “related” whenever a glyph name changes for any reason. Related-glyph renaming has wacky consequences when one glyph name is a substring of another, and the Beikaitoru glyph-naming scripts (among other things in Tsukurimashou) have to take it into account. This issue was reported as FontForge Github issue number 523 (<https://github.com/fontforge/fontforge/issues/523>) and patched in their mainline; I don’t understand FontForge’s release schedule if they have one, but presumably, distributed versions will someday contain the fix. Until we’re sure that day has passed, Tsukurimashou will continue to work around the issue.
- ☹ FontForge “Metrics” window crashes when OpenType feature choice is changed, in the presence of a complicated multi-lookup GSUB feature. Reported as FontForge Github issue number 1138 (<https://github.com/fontforge/fontforge/issues/1138>), January 2014.
- ☹ FontForge crashes when generating BDF if the copyright string is a null pointer, because of using `==` instead of `!=` when attempting to test for that case. Pull request to fix the typo created and merged (<https://github.com/fontforge/fontforge/pull/1143>), January 2014.

X₃T_EX X₃T_EX

- ☹ X₃T_EX fails to advance the glyph pointer after a successful match in a GSUB table, which has complicated consequences for chaining substitutions, most notably that “ignore sub” rules have no effect. This is actually a bug in the third-party ICU library which X₃T_EX uses. Reported to X₃T_EX mailing list in November 2011; after discussion there (<http://tug.org/pipermail/xetex/2011-November/022298.html>) I found an existing issue in ICU’s bug tracker dating from June 2010 (<http://bugs.icu-project.org/trac/ticket/7753>). Poor workaround is to carefully write all substitution features to

work regardless of whether the pointer advances; one useful technique is to make what would be an “ignore sub” rule instead substitute to a series of identical-looking glyphs that are never matched on input to subsequent rules. As of mid-2012 there was interest expressed by XeTeX developers in replacing ICU with HarfBuzz, which is known to handle this case correctly and may have other advantages too.

- XeTeX fontspec package: in some versions has trouble with treating its WordSpace configuration option as a multiplier when font size changes, in a way that is most noticeable when using monospaced OTF fonts; complicated by the fact that it’s hard to define just what should count as a “monospace” font. Discussed at length on the XeTeX mailing list in February 2011 (<http://tug.org/pipermail/xetex/2011-February/020065.html>) and eventually resulted in three items and planned fixes in the fontspec issue tracker (<https://github.com/wspr/fontspec/issues/97>, <https://github.com/wspr/fontspec/issues/98>, <https://github.com/wspr/fontspec/issues/99>). The main bug has now been fixed, and it has also become less relevant to Tsukurimashou documentation since the introduction of the proportionally spaced fonts, but still bears some watching. If they implement all my suggestions, then the current PunctuationSpace multipliers in Tsukurimashou documentation will become much too large and need to be reduced. Earlier versions of Tsukurimashou toyed with workarounds based on the everysel package or on poking into the internals.
- XeTeX and fontspec do not apply by default some OpenType features that are supposed to be applied by default (in particular, “ljmo” and “vjmo”). Workaround: manually request them with the RawFeature option. It is possible that more careful language tagging (specifically of Korean text) in the TeX input might reduce or eliminate this issue. Less relevant now (January 2014) that Mandeubsida no longer uses “ljmo” and “vjmo.”
- Version 2.4 of the fontspec package (which I first encountered in TeXLive 2014) changes the syntax of the `\fontspec` macro to place its optional argument after, not before, the required argument. It claims to be backward compatible, but this is a problem

when the macro is followed immediately by text intended to be written in literal square brackets; as a result, Tsukurimashou user manuals prior to version 0.9 cannot be typeset with new fontspec. Switching documents to the new syntax would break compatibility with earlier versions of fontspec. Instead, recent Tsukurimashou documentation inserts `\relax` commands to prevent square-bracketed text from being interpreted as an optional argument.

外のソフトウェア Other software

- ☹ \LaTeX tocloft package: as of May 2011, it sets an entry without a dot leader by actually requesting an entry with a leader, but with a font-dependent invalid spacing between dots. The DVI renderer is supposed to reject the bad spacing and not set a leader at all. With default font sizes, in most DVI renderers, that results in either the correct appearance (no leader) or an almost-correct appearance (a single dot instead of a leader). But with the larger sizes of the extsizes classes, in some renderers, the resulting DVI file is so invalid as to cause the renderer to blank the rest of the page. Reported to Will Robertson, package maintainer; he acknowledges it, and says will fix (by not requesting dot leaders when they're not wanted) in next version. A workaround involving poking into the package internals to set a smaller invalid spacing is implemented near the top of `doc/bkstyle.tex`.
- ☹ PGF (lower layer of TikZ) shapes.callouts library: in version 2.10 only, ellipse callouts just don't work (causing fatal \TeX errors if attempted), apparently because a macro name was changed in the PGF development process and the change was not propagated to all files. Workaround implemented in our build system consists of detecting version 2.10 and making a local modified copy of the buggy file with the macro name corrected. The problem was apparently known to the maintainers and already fixed in their source code repository before I noticed it. It presumably will not be an issue in any future releases, but buggy 2.10 is the latest

release and widely used as of late 2011; it will probably remain in the wild for a long time.

- The ttx utility of TTX/FontTools will emit an empty Format 2 cmap subtable (containing no code mappings at all) if given an XML file that describes a Format 2 cmap subtable in which all the character codes are single-byte codes. Workaround is to include at least one double-byte code. Reported to their bug tracker (<https://sourceforge.net/p/fonttools/bugs/53/>), January 2014.
- HarfBuzz hangul “shaper” deliberately does not handle the case of multiple conjoining jamo in the same category that could better be expressed as a single cluster: for instance, the case of U+1100 U+1100 (kiyeok kiyek) instead of U+1101 (ssangkiyeok). These two sequences are not canonically equivalent in Unicode, and apparently the longer one is forbidden by the Korean standard KS X 1026-1; nonetheless, Mandeubsida’s GSUB features were intended to handle them both nicely. Worked around this and various other issues (some of which were probably bugs in Mandeubsida) by redesigning the GSUB features, January 2014. Thanks to Dohyun Kim and Jonathan Kew for discussion of these points.

『作りましょう』の使い方

Using Tsukurimashou

Fonts appear in OpenType format in the `otf/` subdirectory of the package. If you just unpack a distribution, there will be four fonts there corresponding to the monospace and proportional versions of Tsukurimashou Kaku and Mincho. If you run the build process, that will create more. For a quick start, all you need to do is install the font files in whatever way is standard for installing OpenType fonts on your system. It is safe to delete everything else in the package, though you might want to keep the PDF documentation files.

Samples of what the different styles look like are in the file `demo.pdf`, which see. Here's a brief summary:

- ④ 作りましょう角 Tsukurimashou Kaku (“Square Gothic”): sans-serif with squared stroke-ends.
- ④ 作りましょう丸 Tsukurimashou Maru (“Round Gothic”): sans-serif with rounded stroke-ends.
- ④ **作りましょうアンビルの的 Tsukurimashou Anbiruteki (“Anvilicious”): extra-bold, rounded sans-serif.**
- ④ 作りましょう天使の髪 Tsukurimashou Tenshi no Kami (“Angel Hair”): very thin hairline display font.
- ④ 作りましょう僕っ娘 Tsukurimashou Bokukko (“Tomboy”): felt marker style.
- ④ 作りましょう明朝 Tsukurimashou Mincho (“Ming Dynasty”): modern serif.

The remainder of this document is reference material describing the special features of the fonts, as well as instructions on how to build and customize your own font files, including the styles that aren't distributed in precompiled form.

OpenTypeのフィーチャー

OpenType Features

OpenType contains a mechanism for defining what are called “features” in a specific technical sense rather than the more general usage of that word. They are identified by four-letter tags, and generally correspond to extra data added to the font that compatible renderers can interpret to provide special typesetting effects. The Tsukurimashou build system’s configure script accepts an “--enable-ot-features=” argument which can be given a list of feature tags, or “all”, each optionally prefaced by an exclamation point to negate it. These are processed left to right, so that a setting like “--enable-ot-features=all,!ccmp” will enable everything except the “ccmp” feature. The default value is “all”.

Be aware that just because you selected a feature during config does not mean you can actually use it. Features are only included in fonts if they make sense for the particular font being built (thus, monospace fonts will not contain proportional-only features) and if all necessary character codes are available (thus, fonts with Japanese glyph coverage will not contain shaping features specific to Korean script). Also, build-time configuration only controls what will be included in the font. Given that the feature is in the font, your renderer must then support the feature you want to use, and you may need to adjust the renderer configuration to tell it to use the feature. Some renderers make that easier or harder to do than others, and some renderers do not turn on by default certain features that I recommend and the standards say should be turned on by default.

Note that although OpenType would permit such a thing, features in any single Tsukurimashou font are NOT language- or script-specific; for instance, the set of features available for Latin script is the same as the set of features available for Japanese script in the same font, and the behaviour of each feature does not change depending on whether the text is tagged as “English” or “Japanese.” This even extends as far as, for instance, the fact that you can really kern kanji in Tsukurimashou proportional fonts. A feature applied to a code point sequence for which it has no meaning (such as “lead jamo shaping” applied to English text)

will simply have no effect. Specific languages are nonetheless mentioned in the font tables for the benefit of FontForge, which doesn't seem to handle defaults very well, and any other software that behaves like FontForge.

Some of the features mentioned in this section are not actually OpenType features in the technical sense; they are instead options that configure the Tsukurimashou build system to control things that should or shouldn't appear in the OpenType font file output. They are given customized four-letter tags and configured through the same interface as true OpenType features because they are somewhat related and it's more convenient than having to have some other interface just for them.

外に分すう Alternate Fractions (afrc)

Tsukurimashou does not contain any special support for diagonal fractions, but it does support vertical or “nut” fractions using the OpenType feature “afrc.” With the afrc feature turned on, any sequence of up to four digits, a slash, and up to four more digits becomes a vertical fraction:

$1/2 \rightarrow \frac{1}{2}$ $34/56 \rightarrow \frac{34}{56}$ $789/123 \rightarrow \frac{789}{123}$ $4567/8901 \rightarrow \frac{4567}{8901}$

This feature works with both the “narrow” digits and slash (ASCII code 47–57, Unicode U+002F–U+0039) and the “wide” ones (Unicode U+FF0F–U+FF19). If the input digits are narrow, and the fraction is one digit over one digit, then the resulting glyph will be narrow (the same width as monospaced Latin characters). Otherwise—with wide input or more than one digit in the numerator or the denominator—the fraction will be the width of a wide (ideographic) character.

$1/2$ [→] $\frac{1}{2}$ [] $1 / 2$ [→] $\frac{1}{2}$ [] $34/56$ [→] $\frac{34}{56}$ [] $3 4 / 5 6$ [→] $\frac{34}{56}$ []

As of version 0.5, this feature works in both the monospace and proportional fonts. However, in order to make it work in proportional fonts

it was necessary to exclude the glyphs involved from the kerning table; as a result, only precomposed fractions will benefit from kerning.

キャプがスモールになって Capitals to Small Caps (c2sc)

This feature replaces capital letters with small caps, as might be desired for typesetting all-caps abbreviations like TEPCO and IAEA. It works much like the “smcp” feature. similar limitations apply.

文脈代替字 Contextual Alternates (calt)

The “calt” feature in general replaces glyphs with other glyphs that look better in certain contexts. In the Tsukurimashou family, it is used only by monospace fonts, to replace Latin capital letters with small caps when they are followed by combining accents. The accents (which do not shift vertically) would otherwise usually collide with the upper part of the letters. Accents in proportionally-spaced fonts are positioned using the GPOS table and will shift vertically when applied to capital letters. Small caps in monospace Tsukurimashou fonts are the same height as lowercase letters (a little shorter than small caps in proportionally spaced fonts) to make this substitution work nicely.

When the “ccmp” feature is enabled, as is default, precomposed accented capitals (which are taller than small cap plus combining accent combinations) will be used wherever possible, and the “calt” feature will only affect unusual combinations not covered by precomposed glyphs.

グリフの併合と分解 Glyph (De)Composition (ccmp)

This feature’s general function is to join and split glyphs that can be thought of as combinations of other glyphs. Exactly what that means depends on the script and the particular font. Often, the results of “ccmp” are needed as input for other features. It is recommended that this feature should be turned on by default in fonts that have it, and that is required by Microsoft’s spec.

In all fonts, this feature will substitute precomposed accented Latin

letter glyphs for combinations of letter plus combining accent, whenever precomposed glyphs exist. The precomposed glyphs generally look better because the accent in the precomposed glyph may change shape or size to work better with its base letter; this issue is especially significant in the case of monospace fonts and capital letters, because the generic accent would usually be too low and is not subject to GPOS positioning. When no precomposed glyph exists, rendering will fall back on the unaccented letter and a generic accent, with the capital-letter issue in monospace fonts partially mitigated by the “calt” feature’s substitutions, if those are enabled.

Similarly, this feature will substitute precomposed voiced kana glyphs for combinations of unvoiced kana plus combining dakuten or handakuten.

In the Mandeubside fonts, this feature does extensive processing of precomposed syllables, conjoining jamo, and jamo clusters, to render each syllable into one glyph for the lead (even if it contains more than one jamo), one glyph for the vowel, and zero or one glyphs for the tail, with the lead and vowel glyphs selected from a choice of alternate forms to take into account the layout of the syllable, which is driven by the vowel and the presence or absence of a tail. In the case of syllables that do not have standard whole-syllable code points (i.e. those containing archaic or uncommon jamo or jamo combinations), the output of `ccmp` is the best we can do. Main-line standard syllables, which have their own code points and corresponding precomposed glyphs, will get those glyphs substituted in when the “liga” feature runs.

合字 Ligatures (liga)

The Tsukurimashou fonts do not contain ligatures for the Latin script at present. Their basic letter forms were designed with monospace setting in mind, where ligatures don’t make sense; and given the importance of keeping everything legible to readers whose native language is Japanese and who may not be familiar with some of the advanced aspects of the Latin script, it was a design decision to make the letters look good without ligatures even in proportional fonts.

The Mandeubside fonts, however, contain an extensive ligature table for hangul, and use of this feature is required to get full support for hangul script. Without it, precomposed syllable glyphs will not be used and every syllable will be rendered in one of the six generic layouts. See the section on Korean language support.

文字に付け方 Mark to Base Positioning (mark)

This feature only exists in the proportionally spaced fonts; a similar effect is achieved in monospace fonts via the “calt” feature and zero-width glyphs. It should be turned on by default. The feature stores data for OpenType renderers to attach accents to letters, allowing for combinations of letter and accent that do not have precomposed glyphs of their own. For instance, U+0071 U+0306 will generate “q̇,” a lowercase q with a breve, which does not correspond to any Unicode code point and would not otherwise be available.

As of version 0.6, support for this feature is limited. Not every base letter you might want will necessarily be available; all the most popular combining diacritical marks exist, but it’s easy to imagine others you might want that are not included; in many cases (especially when capital letters are involved) the mark-to-base version of a combined character ends up looking significantly different from the precomposed version; some of the combinations just don’t look very good; and intended extensions of the scheme to, for instance, allow adding dakuten to Japanese kana don’t exist yet at all. But many common cases should be covered.

点に付け方 Mark to Mark Positioning (mkmk)

This feature works in combination with the “mark” feature above, and it only makes sense to use when “mark” is turned on; like “mark” it is only available in proportional fonts. It stores data used by OpenType renderers to add more marks to glyphs that already have some marks; in particular, this should allow stacking up more than one accent on the same letter to create many glyphs that would be impossible by other means.

In version 0.6, this feature has been expanded to cover a few more of the Latin accents than before, allowing the construction of such characters as ě. It is still somewhat experimental, however.

メタデータ Metadata Table (name)

This is not the usual kind of OpenType feature, but it can be turned on and off through the same build-system interface. By default, the font will contain a “name” table listing metadata such as creatorship and licensing, in both English and, as appropriate, one of either Japanese or Korean. Some of this information is obligatory, so if you turn off the “name” feature you actually still get a name table, but it will be less completely populated.

巴の花形 Ornaments (ornm)

Tsukurimashou provides eight tomoe ornaments that look like this:



These glyphs are encoded to the Unicode private-use code points U+F1731 to U+F1738, and always available that way. If the “ornm” feature is turned on, then they will also appear as substitutions (not alternates!) for the ASCII capital letters A through H; so that the text “A BIG TEST” comes out as “● ◉ ◊ ○ ◌ ◍ ◎ ●.” It is likely that the way this OpenType feature works will change in the future, since it seems not to be current best practice to implement ornm by substitution but I’m not quite sure yet what the best practice actually is.

These ornaments look the same in all the fonts; they do not change from one style to the next.

FontAnvilだけのメタデータ FontAnvil-specific Metadata (pfed)

This is not the usual kind of OpenType feature, but it can be turned on and off through the same build-system interface. When selected, it causes the build system to add a “pfed” table with a “flog” sub-

table in the generated OpenType fonts, containing verbose information about the installation on which the fonts were built. This table is a FontForge/FontAnvil-specific extension of OpenType format; the name “pfd” refers to PfaEdit, an earlier name of the software that later became FontForge. Other software will ignore this table; but users with FontForge can examine the verbose metadata under the “FONTLOG” heading in the “Font Info...” dialog. The data is also visible as a chunk of plain text near the end of the file when the OTF file is examined with a general-purpose file viewer such as “less.” The main interesting content is the command line that was given to configure, and a dump of most of the variables known to Make. There’s also a copy of the copyright notice and URLs for the project home pages, giving in more verbose detail some of the same information included in the “name” table.

I recommend activating this feature, especially if you will be building fonts for distribution to others. It may make bug reporting easier, because it means that anyone who gets a copy of the font also gets some information on where that font came from; if someone builds a font and then has trouble with it, it’ll be easier to help if the font contains this debugging information. It also improves the chances that should a font file get separated from its context, someone stumbling upon it will be able to figure out what to do with it. The Net is full of inaccurately labelled fonts with unknown authorship, often being sold by shady commercial enterprises that have no legal right to do so, and we all ought to do what we can to stamp that out. Detailed human-readable metadata, in general, is a Good Thing.

But the font log contains information like software version numbers, user account names, and installation directory names. Some people have funny ideas about the sensitivity of such information in relation to system security; they may think that revealing it creates a real risk, not otherwise present, of people breaking into their computers. They might also think it could be used to trace the origins of anonymously written PDF files and that forensic investigators don’t have many other ways to do that.

Such people are wrong. However, my saying so won’t cause them

to change their minds. If I distributed software that attached this information to generated fonts by default, then someone who didn't read the documentation would eventually "discover" it and make a big fuss about it supposedly being a security hole. Who needs that? The feature is therefore turned off by default. I recommend turning it on by passing the "--enable-ot-features=all" option to configure. The default is "--enable-ot-features=all,!pfed," which enables everything else.

Regardless of the setting chosen, the build system will place the same information in a file called "fontlog.txt" in the txt/ subdirectory of the build tree. After doing a build you can read that file to see what would have been put in the fonts. The option setting just controls whether or not the fontlog file will be added to the OTFs during the final packaging step.

スモールキャピタル Small Caps (smcp)

SMALL CAPS ARE AVAILABLE THROUGH THE OPENTYPE "SMCP" FEATURE. This feature simply substitutes the 26 lowercase ASCII Latin letters with alternates that are encoded into the Private Use Area at the code points formerly used by Adobe for this purpose: U+F761 to U+F77A. At some point in the future, other glyphs may be added to support small cap letters other than the 26 ASCII ones, and I may start using unencoded glyphs rather than mapping them into the PUA.

In monospace fonts, these small caps are about the same height as most lowercase letters, in order to allow their use as substitutes when combining accents are applied to uppercase letters. In proportionally spaced fonts, small caps are a little taller than lowercase letters, but shorter than ordinary caps.

ヘビーメタルウムラウト Heavy Metal Umlaut (ss01)

With Stylistic Set 1 (OpenType feature "ss01") turned on, umlaut or dieresis over the vowels ÄËÏÖÜÿäëïöüÿ and tilde over Ñ and ñ are replaced by a "heavy metal" umlaut intended for spelling musical names like "Motörhead," "Spiñal Tap," and "Mormoiñ Tabärnacle Choïr." The

heavy metal umlaut differs from the regular umlaut in that the dots are larger and pointier.

Glyphs to support this feature, including a “spacing heavy metal umlaut” character, are encoded into the private-use area at U+F1740 through U+F174E.

丸つき字 Enclosed Letters and Numerals (ss02)

With Stylistic Set 2 (OpenType feature “ss02”) turned on, the enclosed characters described by Unicode become available as contextual substitutions for sequences of (in most cases ASCII) characters:

- ⓪ (0) → ① through (50) → ⑵
- Ⓐ (A) → Ⓐ through (Z) → Ⓔ
- ⓐ (a) → ⓐ through (z) → ⓷
- Ⓐ (ア) → Ⓐ through (ン) → ン
- Ⓐ ((1)) → ① through ((10)) → ⑩
- Ⓐ {0} → ① through {20} → ⑳
- Ⓐ {A} → Ⓐ through {Z} → Ⓔ
- Ⓐ [A] → Ⓐ through [Z] → Ⓔ
- Ⓐ <A> → Ⓐ through <Z> → Ⓔ

The choice of which ranges of numbers and letters to support is mostly not mine—I just implemented what I found in the Unicode charts.

Unicode only has code points for the unvoiced versions of the enclosed katakana (e.g. ⓐⓑⓒ but not enclosed versions of ガギグ), and it has no code point for enclosed ン. I’ve added ン, encoded to private-use code point U+F1711, but not the others. If there’s a demand for other

enclosed characters, though, they are pretty easy to add. I held off on just defining dozens or hundreds more, because it's not clear to me what the typical use of these characters actually is, and a complete set of enclosed characters might be better seen as a font in itself rather than a series of special characters inside the font.

Note that for the enclosed katakana the substitution will accept either ASCII parentheses or wide parentheses (U+FF08 and U+FF09); the others work with ASCII parentheses only.

私用符号位置を避けて Limit PUA Code Points (xpua)

This is not the usual kind of OpenType feature, but is configured through the same interface as the others.

For software engineering purposes it is convenient to have a unique Unicode-like code point for every glyph in the font family. However, some of those glyphs are not intended to be accessed directly by users, and more than one of them may correspond to the same Unicode character in ordinary text. This feature removes the private-use code points from many such glyphs.

For instance, in Mandeubside there is a glyph named “uni112A” for code point U+112A, but there is also a glyph named “uni112A.bug” which looks identical but is referenced in the substitution table in order to work around a bug in the ICU library used by Xe_{La}TeX. The uni112A.bug glyph, as its name implies, is used to typeset the character U+112A, but it is referred to internally using the code point U+FF72A. That is a code point in one of the Unicode Private Use Areas (PUAs), which are special ranges of code points reserved by Unicode for purposes like this. The U+FF72A code point is the one that could, for instance, be used to enable or disable the uni112A.bug glyph in the `--enable-chars` option to the `configure` script.

The Tsukurimashou package uses hundreds of these PUA code points. They are necessary for keeping track of things during the compilation and assembly of the fonts. However, most of the glyphs with PUA

code points are not intended to be directly accessed through their code points, only indirectly through glyph substitution. Having an unnecessary PUA mapping in the font for any given glyph may cause some buggy software to ignore the glyph's name, resulting, for instance, in incorrect text when cutting and pasting from a PDF file. Encoding glyphs to PUA code points could also possibly conflict with other uses of the same PUA code points by other software. So all in all it is usually preferable that these glyphs should not have any code point mappings at all in the finished fonts. On the other hand, removing their code point mappings means it is no longer possible to access the glyphs directly by code point; then short of using third-party software to bypass the mapping tables and access glyphs by raw glyph numbers, users **must** use the glyph substitution features instead, which means both that the glyphs are only available in the contexts where the substitution features are designed to provide them, and that the users must have OpenType rendering software that at least sort of works (which is not necessarily common in the field).

The “xpuā” feature allows a configuration choice between these two situations. With it turned on, which is the default, most of Tsukurimashou's PUA glyphs will have their code points stripped in a late stage of font packaging. It is not universal: some glyphs that do not correspond directly to any Unicode character remain accessible at PUA code points because that's relatively harmless and they may be difficult to invoke in other ways. But glyphs that exist only to support OpenType features (such as the ones needed for vertical fractions and Korean syllable layout) end up as unencoded glyphs in the final fonts. As of version 0.7, the Tsukurimashou PUA-encoded characters that keep their encodings in the final fonts notwithstanding “xpuā” are exactly the glyphs in the U+F17xx range.

Turning off “xpuā” causes every glyph to be encoded at its internal code point. That was the state of affairs in versions of Tsukurimashou prior to 0.7, which didn't offer a choice.

付加文字

Extra glyph coverage

げんじもん Genjimon

Glyphs for the 54 Genjimon are encoded to private-use code points U+F17C1 to U+F17F6, in the order of the corresponding Tale of Genji chapters (“Kiritsubo” to “Yume no Ukihashi”). The style of these glyphs varies a fair bit between the different font styles. Here are samples:


- ☪ Kaku: 
- ☪ Maru: 
- ☪ **Anbiruteki:** 
- ☪ Tenshi no Kami: 
- ☪ Bokukko: 
- ☪ Mincho: 

A set of Genjimon-only fonts, derived from the Tsukurimashou code, exists as a spin-off project in the genjimon subdirectory of the Tsukurimashou distribution. The documents there contain some more information about the history of these symbols and what they signify.

えききょう I Ching

All the I Ching-related characters defined by Unicode in the Basic Multilingual Plane are supported: trigrams U+2630 to U+2637, monograms U+268A and U+268B, digrams U+268C to U+268F, and hexagrams U+4DC0 to U+4DFF. These characters are sized to fit in the same box as kanji, and so in most cases you will probably want to scale them to a significantly larger point size than that of nearby text. The style does change somewhat from one font to the next.

There are also some “I Ching dot” combining characters in the private-use code points U+F1701 to U+F1709, intended to show movement of lines 1 through 6 of hexagrams or 1 through 3 of trigrams, encoded in that order. The idea is that you can typeset something like this:



using a sequence of codes like this:


U+4DDE U+F1702 U+F1704 U+2192 U+4DEF

解字の文字 Ideographic Description Characters

Unicode defines a set of special code points (U+2FF0 to U+2FFB) and a syntax for using them to describe what Han-script characters look like. For instance, 「齒」 might be described by the string 「☐止☐☐米」, which expresses that 齒 consists of 止 written above a composite sub-glyph that itself consists of ☐ wrapped around 米. Unicode’s scheme for character description appears to have been inherited more or less unmodified from a similar scheme in the Chinese GBK standard.

Tsukurimashou includes glyphs for the special characters. There is also support in the build system for a “make eids” target, which generates a file called tsukurimashou.eids in the txt/ directory. That file is a collection of Extended Ideographic Description Sequences (EIDSes) describing the construction of the kanji in the font. This is the interface to IDSgrep, a program for searching kanji according to powerful criteria on their visual structure. IDSgrep is available from the same OSDN Japan project that hosts Tsukurimashou; see that package and its documentation for more information on EIDS format and how to use it.

With or without IDSgrep, Tsukurimashou may be used to typeset the special characters for the descriptions. The glyphs look like this:



公のユーログリフ Official Euro Sign

Fonts that include a symbol for the European currency unit (Unicode U+20AC) usually re-draw it to match the style of the rest of the font, and the Tsukurimashou fonts are no exceptions. However, there is an official glyph design that apparently was intended to be normative—in theory, you’re supposed to use the official glyph, which does not change with the style of the rest of the font, even if it clashes. For those who want to do so, the official Euro sign is included in the Tsukurimashou fonts at private-use code point U+F1710.

It looks like this: €

ハングル語

Korean Language Support

This user manual was built without the Korean fonts (in particular, Mandeubsida Dodum PS) and as a result, the section on Korean support cannot be typeset.

フォントの名前について

Regarding Font Names

First, I will use the term “font” to refer to an OTF or similar file with a complete glyph set (or as many glyphs as have been selected with configure). This will generally be a few thousand glyphs for the Tsukurimashou family and about 13 thousand for the Mandeubsida family. Because METATYPE1 is limited to producing Postscript font files with at most 256 glyphs each, the system builds “subfonts” each corresponding to a “page” of the Unicode character set—that is, an aligned 256-code-point block. For instance, the range U+0000 to U+00FF is the page containing ASCII and Latin-1 characters.

The division into pages provides some advantages for the build system: when something is changed on one page it may be possible to avoid re-compiling the others, and if there are many pages to be compiled then they can be compiled simultaneously on a multi-core computer. So we would probably choose to keep this design feature even if some future version of the system were no longer bound by the Postscript 8-bit limit that originally made it necessary.

The most authoritative form of the name of a font is the “Hamlog name,” which is an ordered quadruple of Hamlog atoms. (The Hamlog language is described elsewhere in this document.) Hamlog names are manipulated by, for instance, the code in `select.hl`. The items in the quadruple are called the family, the style, the weight, and the spacing. An example Hamlog name might be “(tsukurimashou, bokukko, demibold, monospace).”

The family represents the broad category of the font. Generally, all fonts sharing the same family value will have basically the same glyph set and the same overall shapes for all glyphs, though the visual style of the strokes will vary. If there are significant differences in shapes or glyph sets, they will be put in a separate family. At present the supported values for the family are “tsukurimashou,” “tsuita,” and “mandeubsida,” though there is also some experimental (less than alpha status) code for families to be called “blackletter_lolita” and “kazoe-

mashou.”

The possible values for style will depend on the family. Styles generally correspond to sets of preset values for the adjustable parameters of a family’s letter-drawing code. The style values currently allowed, per family, are as follows. Style names marked with an asterisk are not allowed as part of the basic Hamlog name but may be introduced by rewriting, as explained below.

tsukurimashou kaku, maru, mincho, bokukko, anbiruteki*, tenshinokami*.

tsuita soku, atama.

mandeubside dodum, batang, sunmoon.

blackletter_lolita cosette.

kazoemashou seisuu.

The weight may be “extralight,” “light,” “normal,” “demibold,” “bold,” or “extrabold.” The spacing may be “monospace” or “proportional.” All combinations of these are allowed for all families except that when the family is “kazoemashou,” the weight and spacing must be “normal” and “proportional” respectively; that constraint is enforced by the “allowed_style/4” predicate in select.hl.

The Hamlog name may be subject to rewriting according to rules in the select.hl file, specifically the “do_style_xlat/8” predicate. These rules support a few style/weight combinations that have been given unique style names of their own, both because that’s fun and as a legacy of earlier versions in which the weight parameter was not fully implemented. This code is also meant to support future per-language localization of name components like “demibold,” in particular so that Korean fonts won’t end up with Japanese in their names, though that aspect is not yet fully implemented. Here are the rewriting rules:

- ⊗ (tsukurimashou, maru, extralight, *X*) is rewritten to (tsukurimashou, tenshinokami, normal, *X*).

- ④ (tsukurimashou, maru, extrabold, X) is rewritten to (tsukurimashou, anbiruteki, normal, X).
- ④ (kazoemashou, X , normal, proportional) is rewritten to (kazoemashou, X , normal, monospace). Bearing in mind that “monospace” translates to the empty string in the long name, the point here is to avoid a redundant “PS” in the long name (see below) for the Kazoemashou family, where there is only a proportional version.

The Hamlog name is translated into a few other kinds of names, by a combination of Hamlog code in `select.hl` and Perl code embedded in `Makefile.am`. One important translated name is the “short name,” which is a sequence of two, three, or four short ASCII tokens, which in different parts of the code may be separated by either hyphens or underscores. An example short name is “tsuku-mi-ps,” corresponding to the Hamlog name “(tsukurimashou, mincho, normal, proportional).” Each element in the Hamlog name is translated into a token for the short name, according to the table given by the `short_name/2` predicate in `select.hl`, but some Hamlog atoms (namely “normal” and “monospace”) translate to the special value “l___da” (for “lambda,” meaning empty string; triple underscore marks it as requiring special handling in the Perl wrapper) and are removed, often leaving fewer than four tokens in the result. Short names are always derived from the Hamlog name **before** rewriting.

Subfonts also have short names of their own, formed by adding an extra token at the end which is the two- or three-digit hexadecimal representation of the page number. For instance, “tsuku-mi-ps-00” would be the short name of the subfont covering U+0000 to U+00FF of “tsuku-mi-ps”. These names are often used in constructing filenames for intermediate files during build.

Anyone contemplating modification of this system should be aware that short-name tokens other than page numbers must **not** be valid hexadecimal digit sequences that could be mistaken for page numbers; in other words, each non-page-number token should contain at least one letter in the range “g” to “z.” The build system needs to be able to distinguish subfont and full-font short names from each other by

testing for a hexadecimal number as the last token. To prevent terms like “extra bold” and “demibold” from abbreviating to valid hexadecimal numbers, terms derived from the word “bold” are abbreviated as if it were spelled “qold” or “bqld.”

Fragments of short names are also used to name source code files. Files that follow this scheme (not all source files do) have names consisting of one family token, a hyphen, and one other token, followed by the appropriate extension. Building a subfont involves creating a temporary driver file that includes all the relevant source files for that font’s short name. For instance, to build “tsuku-mi-ps-00.pfb” the build system would create and execute a driver file that loads “tsuku-mi.mp,” “tsuku-ps.mp,” and “tsuku-00.mp,” in that order. Some of those files might themselves load other files outside this naming scheme. Almost all file inclusion is done via conditionals in METAFONT to prevent any file from being included more than once, but double inclusion should also be both avoided by design and usually harmless should it occur by accident. The file inclusion macros in preintro.mp also support a queue of “late” includes, which will be processed after the others to allow (for instance) code associated with a style to override code associated with a weight.

In the case of the family short name token not being “tsuku,” the build system will use source files that match the subfont’s family token where those exist, and default to “tsuku” when such files don’t exist. For instance, to build “bll-co-01.pfb” (Blackletter Lolita Cosette, page 01) it might load “bll-co.mp” (which exists), but then load “tsuku-01.mp” because “bll-01.mp” doesn’t exist. Thus, the “tsukurimashou” family can be thought of as a generic ancestor from which any others inherit.

Fonts also have “long names,” which are created by a similar translation process. Long names are always derived from the Hamlog name **after** rewriting. Long names are generally mixed-case ASCII, and the elements are separated by spaces and may also contain internal spaces. As with short names, normal weight and monospace translate to nothing and are removed when forming the long name. An example long name might be “Tsukurimashou Tenshi no Kami PS”; note that that would correspond to the short name “tsuku-mg-el-ps” (as if “Tsukurimashou Maru

Extra Light PS”) because of rewriting. There are also some places, notably in the filenames of final OTF fonts, where a long name is used with all the spaces and hyphens removed.

There are also “JK names,” which are constructed just like long names but with elements written in Japanese or Korean, and thus generally in non-ASCII Unicode characters. For Mandeubsida fonts these should be in Korean (but the occasional Japanese or English word may appear because of a yet-unimplemented translation feature) and have spaces between the elements; otherwise JK names are in Japanese and have no spaces. The main use of JK names is for populating the font metadata tables. An example Japanese name might be 「作りましょう天使の髪PS」, corresponding to “Tsukurimashou Tenshi no Kami PS.”

Since there is some information loss in transliteration to English, and many of the names come from Japanese and Korean words anyway, it may be best to think of the JK name as the truest name of the font for human purposes, even though from the code’s perspective the unrewritten Hamlog name is the true name of the font and all others are derived from it. The source of the translation table for JK names is in `txt/jnames.txt`, which is processed by Perl to change the UTF-8 into a hexadecimal form that Hamlog can handle, the result going into `hamlog/jnames.hl`.

『作りましょう』を作りましょう!

Building Tsukurimashou

Please note that if you just unpack the Tsukurimashou distribution and look in the “otf” subdirectory, you will find some ready-made font files there. If you are content with them, then those are the only files you need; you can safely delete everything else in the package (maybe save the PDF documentation files, if you want) and ignore this section. These notes on building Tsukurimashou are only for expert users who want the greater control, wider style coverage, and intellectual challenge of doing custom compilation.

First I will describe building from a source distribution package. If you wish to compile from a version-control checkout, then also read the specific notes about that in Subsection ???. If you want to create distribution packages, follow the procedure described in Subsection .

Assuming you wish to embark on the adventure of compiling Tsukurimashou, you must have at least the following:

- ④ A reasonably standard Unix command-line environment. I use Slackware Linux. Anything branded as “Unix” should work. MacOS X might work. Windows with Cygwin might work.
- ④ A standard C-language tool-chain (normally comes included with Linux).
- ④ Perl.
- ④ GNU Make 3.82 or later (non-GNU versions will not work; earlier versions will not work).
- ④ Metapost (this comes with most T_EX/L_AT_EX installations)
- ④ An interpreter of PE scripts. FontAnvil is preferred, and is included with the Tsukurimashou sources. FontForge will also work, as long as it retains support for “native” scripting.

Other things that might also be useful include:

- ④ $X_{\text{Y}}\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ (needed for compiling the documentation).
- ④ Expect (makes the build system work better—exactly why is explained in more detail in the warning message you will get if you don't have it).
- ④ The KANJIDIC2 database (needed for the kanji coverage chart and a planned future fine-grained character subset selection feature).
- ④ A supported Prolog interpreter (not necessary, but if one is available then the build system will use it and run more efficiently).
- ④ A multi-CPU computer (the build system will by default detect and use all your CPUs; since some stages of compilation require a lot of processing, it's nice to have several).
- ④ A checksum program, preferably sha256sum (used for some subtle build-system optimizations).

If you have a current version of $\text{T}_{\text{E}}\text{X}$ Live, then you probably have Metapost and $X_{\text{Y}}\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ already. Note that it does have to be fairly up-to-date. Tsukurimashou now bundles its own version of the public domain METATYPE1 code, so the dependency of previous versions on the mtype13 distribution of that has been removed from the list. In preparing METATYPE1 code for bundling I discovered a previously unknown dependency of previous versions on the t1asm package from t1utils; that has now been bundled as well.

Tsukurimashou compilation depends heavily on scripts written in the FontForge “native” scripting language, which I call PE script for neutrality. Although (as of July 2014) mainline FontForge continues to be able to interpret this language, the FontForge development team treats PE script as a legacy feature and encourages users to switch to Python instead. It is highly unlikely that the Tsukurimashou project will ever switch to Python. To ensure continued support for Tsukurimashou, I now maintain a standalone PE script interpreter called FontAnvil; this should

be available wherever you got the Tsukurimashou source code, and is the preferred PE script interpreter for building Tsukurimashou. As of version 0.11, both FontAnvil and FontForge will work, more or less equally well. Continued FontForge compatibility in the future is not guaranteed.

ビルドのシステム Build system

The build system is based on GNU Autotools and should be reasonably familiar to anyone who has compiled popular free software before. Run “./configure --help” for a description of available configuration options; run “./configure” possibly with other options as described in the help message to set up the system the way you want it; then once it is configured to your liking, run “make” to build it. The completed font files end up in a directory called “otf” in the distribution tree. A few ready-made ones should be there already when you unpack the distribution, for the benefit of the probable majority of users who can't do their own compilation.

If you run “make install” it will attempt to install the fonts and documentation files in sensible locations, but it's not really customary for fonts to be installed like other software, and you may be better off simply taking the compiled files from the otf directory and doing with them whatever you would normally do with fonts instead of using the automated install. At this time there is no T_EX-specific font installation support, though that might be a nice feature for me to build in the future. The install target uses the “--prefix” and similar options to configure in a reasonably obvious and standard way.

If you have KANJIDIC2, place it (in the gzipped XML form in which it is distributed) in the build directory, the doc subdirectory off of the build directory, your system's dict or share/dict directories, possibly under /usr or /usr/local or your configured installation prefix, or your home directory or “~/dict”; or put it somewhere else and tell configure where with the “--with-kanjdic2” option. Although what's used by Tsukurimashou is factual information not subject to copyright, and KANJIDIC2 (which also includes creative content that might be

subject to copyright) is distributed freely, it is not distributed under a GPL-compatible license and so for greater certainty I'm avoiding including KANJIDIC2 or any data extracted from it in the Tsukurimashou distribution. What you get if you have KANJIDIC2 is any features that depend on knowing the grade levels and similar properties of kanji characters—such as the chart showing how much of each grade has been covered, which is important for my own development efforts in planning what to design next and knowing when to release.

The build system for Tsukurimashou is fairly elaborate; it may seem like overkill, but given that I'm running it on a daily basis over the course of many years, it's important that it should be bulletproof, pleasant for me, and efficient in time consumption. Thus it defaults to “silent” build rules, uses pretty ANSI colours, and does a bunch of complicated checks on whether file contents (not just modification times) have changed in order to avoid expensive dependency rebuilds unless those are really needed, while triggering them automatically when I add new source files.

One important tip is that if things are failing, you can add “V=1” to the “make” command line, as in “make V=1 pfbtmp/tsuku-kg-00.pfb” to temporarily disable the silent build rules and see what's going on. Note no hyphen, because this is a variable setting instead of an option, and you will probably also want to specify a target filename, which implicitly overrides the default “-j” multi-CPU option. In version 0.6, Expect was rendered less necessary by the replacement of mtype13 with a homegrown and less chatty Perl script; but the support remains in place to handle T_EX's lesser jubilations.

You can also add “KEEPTMP=1” to prevent deletion of the temporary directories created while running Metapost. This feature is primarily useful for debugging the pfb-generating scripts themselves.

Several options to configure share a common argument syntax under which you can specify a comma-separated list of commands to include and exclude subsets of things. This syntax is actually translated into Hamlog code and evaluated by the logic-programming system during build. It may be easier to understand by example than to explain

rigorously, but here is an attempt at a precise explanation.

A specification consists of a sequence of commands separated by commas or spaces. Each command either adds or subtracts some subset of the things being configured, to or from the currently selected subset which is initially empty. A command consists of one or more alphanumeric tags representing subsets, separated (if more than one) by ampersands or colons, which mean “AND,” and optionally prefixed by plus, minus, or exclamation point. If the prefix is “-” or “!” then the AND of all the tags will be **subtracted** from the current selection; otherwise it will be **added**. The allowed values of the tags depend on which option we are considering, but typically include “all” and “none.”

For example, the specification “mushroom&purple” would select purple mushrooms. The specification “fungi;-mushroom&purple” would select all fungi that are not purple mushrooms. The specification “mushroom;-red,tomato” would select all mushrooms that are not red as well as all tomatoes **including red tomatoes**; tags are evaluated left to right and each one overrides the results of all previous tags.

The “--enable-chars” option uses the above syntax to describe which characters (Unicode code points extended by internal private-use allocations for some unencoded glyphs) should be included in the fonts. The default is “all,” meaning all characters defined in the source code will be included in the fonts. Other tags currently supported are “none,” “ascii,” “latin1,” “mes1,” and families of tags named like “page12,” “uni1234,” and “u123ab,” with lowercase hexadecimal digits, which correspond to 256-character blocks of Unicode code points and to individual code points. For example, the specifier “all;-page02” would include every character it can except none in the range U+0200–U+02FF; the specifier “uni0041” would create very small fonts containing only the uppercase A from the ASCII range. The default is “all.” It is planned that in some future version this option should also be able to use KANJIDIC2 data to allow selection of kanji characters by criteria like school grade level, and some support for that exists already, but it probably doesn’t work yet.

The “--enable-ot-features” option works the same way for OpenType features; note that OpenType features may also be automatically and

silently disabled in whole or in part, overriding this setting, if you have used the previous setting to disable characters necessary to implement the features. For instance, you can't have OpenType contextual substitution support of fractions or enclosed numerals if you have disabled the numeral characters—though you **can** build a font with enclosed and not regular numeric glyphs, because glyphs are mostly independent of each other.¹ This automatic disabling may not be perfectly clean, either; in some cases disabling a character might not disable the feature code that mentions it, and in that case FontAnvil may create an empty glyph for the character even though you disabled it. That shouldn't happen unless you attempt something very unusual and non-standard, however. Currently supported tags are “all,” “none,” and a four-character tag corresponding to the OpenType name of each feature that exists, and a few that don't. See Section for a detailed list.

The “--enable-styles” option allows similar selection of type styles, For instance, “tsuku&kg;el” would enable all Tsukurimashou Kaku fonts, in all weights and both monospace and proportional, **except** no “Extra Light” fonts. The atomic tokens for this selection include the standard “all” and “none,” and the short name tokens and Hamlog atoms for all the families, styles, weights, and spacings currently supported. The token “proofs” may also be used as a shortcut for enabling all the styles needed to generate proof documents. See Section for more about the naming of fonts, and note that these are the pre-rewriting names, so that for instance Anbiruteki might be selected via “tsuku&mg&eq.”

The “--enable-parasites” option allows this kind of include/exclude selection for parasite packages (see next subsection). The supported tokens are “all,” “none,” and the names of the parasites, currently “beikaitoru,” “fontanvil,” “genjimon,” “icemap,” “idsgrep,” “kleknev,” and “ocr.”

The default parasite setting is “auto” (note: not a supported token itself); this enables a special bootstrapping mode in which the configure script will conditionally enable fontanvil and icemap as necessary for the local

¹Mostly. In the case of white-on-black reversed glyphs and some fractions precomposed by FontAnvil instead of by METATYPE1, you must include all the parts that FontAnvil will assemble in order to get the combined glyph made by assembling those parts. This is a sufficiently arcane scenario that the build system will not check for it.

configuration. Normally, that means FontAnvil will be built if there is no FontForge or FontAnvil already on the system; Icmmap is not required for building from a source distribution, but will be built in a version control checkout situation.

The “make dist” target defaults to building a ZIP file only, instead of GNU’s recommended tar-gzip. This decision was made in order to be friendlier to Windows users, who tend to have seizures when confronted with other archive formats, and T_EX users, who are accustomed to ZIP as well. However, the makefiles should also support many other formats via “make dist-gzip” and similar.

I am not confident that “make dist” will really include everything it should if you have disabled optional features with the above options to configure. Please “make me one with everything,” as the Buddhist master said at the hamburger stand, before trying to build a distribution.

The “make clean” target and its variations should make things pretty clean, now that we’re attempting to make the system pass “make distcheck” before each release. Use them with caution, however.

There is a “make check” target, which will run FontAnvil’s fontlint program on all the OTF files. This is a very stringent test and it’s quite likely that fonts will fail despite really being in pretty good shape; fontlint flags a lot of things that I am not convinced are actually errors at all; so if you run it and get nothing but complaints, don’t panic. At the very least, you will need a PE script interpreter with my all-extrema patch (it is included in recent versions of FontForge and all versions of FontAnvil) to have much hope of building fonts that can pass fontlint. Running “make check” at the top will also recursively run “make check” on any parasite packages that have been enabled, and the parasite package test suites (in particular, the one for ID_Sgrep) may be of more practical use than the main Tsukurimashou fontlint tests.

If you try to do “make distcheck,” it will fail if any of the tests fail; and the tests probably **will** fail because of the bugs in FontAnvil spline geometry among other reasons. See the notes in Subsection for workarounds.

Some C source code in FontAnvil and IDSGrep needs to be automatically generated by Icmmap, from data files downloaded from the Unicode Consortium. Distribution packages will contain pre-generated versions of these source files, so that neither Icmmap nor a download is needed. However, if you are building from a source control checkout, or possibly if you have Icmmap installed, the build system will attempt to regenerate these files, and in that case, you must run “make download” before running “make” to get the needed input files.

Don't bother with “--disable-dependency-tracking”; that is an Autotools thing meant for much larger and more software packages. It applies only to code in languages supported by Autotools, which at present means only the C kerning program whose dependency is trivial. The dependency tracking for METATYPE1 code is completely separate, unaffected by this option, and trying to disable it would be a bad idea.

Autotools encourages the use of a separate build directory, with the sources remaining inviolate elsewhere, but that is not recommended for Tsukurimashou. I try to make it pass “make distcheck” right before each release, which implies making separate build directories work, but if you are building Tsukurimashou from a checked-out SVN version, it may or may not work. It's safer to build right in the main directory. Even if VPATH builds work, they are only intended for the case of having an untouched set of sources in one directory and a build in another. If you try to do the overlay thing, with modified versions of some source files in your build directory, it is unlikely to work, because of the large amount of bolted-on filename and path logic that doesn't go through GNU Make for name resolution.

If you look in the source of the build system, specifically files like `configure.ac`, you'll see that I did a whole lot of work ripping out sections of Autotools that were designed for installations of executable software. GNU standards require the definition of a ridiculous number of different installation directories, almost none of which are applicable to a package of this type, and I took out most of the support for those to reduce the cognitive load for users who would otherwise have to think about their inapplicability. The main Tsukurimashou package doesn't install any executables, libraries, C header files, or similar, at

all. (Some of the parasites do.) Cross-compilation and executable name munging were removed for the same reason; some C programs are compiled for kerning and Type 1 opcode assembly, but these are only meant to run on the host system during build, with all the installable files being architecture-independent. The hacking I did on Autotools means that if you modify the build system such that you would be re-running Autotools, it's likely to break unless your version of Autotools is close to the version I used. The configure script will try to detect such a situation and warn you.

The design of this build system was influenced by Peter Miller's interesting article "Recursive Make Considered Harmful," available at <http://miller.emu.id.au/pmiller/books/rmch/>. The fonts as such are all built from a single large Makefile. However, in the current version there is also some use of recursive Make and recursive Autotools to support the parasite packages (next section), should those be enabled. Some notes on the design of the build (prior to the introduction of parasites) are in my Web log article at <http://ansuz.sooke.bc.ca/entry/226>. I will probably write a similar article about the parasite mechanism at some point in the future.

寄生パッケージ Parasite packages

Some of the technology developed to support Tsukurimashou has more general applicability. For instance, the stripped-down version of METATYPE1 bundled with Tsukruimashou may be useful for other fonts that build from METAFONT-language source files to produce modern vector file formats. It would be nice to be able to manage these kinds of associated packages in a way that maintains their link to Tsukurimashou—so that for instance updates in the shared code will propagate to all packages that use it—but nonetheless allows the other packages to be distributed as separate entities to users who don't want a copy of the entire Tsukurimashou system. There are relevant side issues, too, such as a hope for attention directed at one package to attract interest to other packages.

One possible approach might be to attempt to factor out code that

will be shared into some kind of “Tsukurimashou library” package that could then be distributed separately and made a dependency of all the other packages that use it. A problem with doing that is that which code is or isn’t shared will vary a lot due to the wide range of different packages that might want to share some part of the system. Instead, the approach introduced in Tsukurimashou 0.7 uses what I call “parasite packages.” These are sub-packages for Autotools purposes.

Each parasite appears in its own subdirectory of the full Tsukurimashou Project distribution tree or SVN checkout. By default, they are distributed in distribution packages of the full project, but not built by a top-level build. They can be activated using the `--enable-parasites` option to the top-level configure script, but each one in its own subdirectory is also a potentially standalone Autotools package, with its own configure script and capable of building a standalone distribution of itself with the “make dist” target. Parasites will generally look for a Tsukurimashou Project build in the parent directory and connect to it if one is found.

The current parasites are Beikaitoru, FontAnvil, Genjimon, Icemap, IDS-grep, Kleknev, and OCR. Each of these has its own detailed documentation, which see; the notes here are just a summary to give some idea of what the packages are all about. Most of these existed in some form prior to the introduction of the parasite build mechanism, and in some cases vestiges still exist of their old, less systematic, ways of connecting to Tsukurimashou. In the future, the connections will be more streamlined and normalized.

Note that parasite packages generally have their own version numbers which are not necessarily synchronized to specific Tsukurimashou releases; in some cases a Tsukurimashou release will contain a pre-release or even a broken version of a parasite. But it should usually be reasonably safe to mix and match parasites and Tsukurimashou versions of reasonably similar vintage. Parasite packages may also have their own license terms, but must always have license terms that allow them to be distributed within the GPL3 top-level project distribution.

Beikaitoru A modern revival of the historic plotter fonts “digitalized” [sic] by A. V. Hershey of the U. S. Naval Weapons Laboratory during the 1960s, and subsequently distributed by the U. S. National Bureau of Standards and on Usenet. This version includes the Japanese character range omitted from most other revivals. GPLv3 with font embedding exception.

FontAnvil A PE script interpreter, for editing fonts under program control. Overall package GPLv3 or later, some parts also available under BSD-style and other licenses.

Genjimon A set of standalone fonts for the Genjimon symbols, which in the main Tsukurimashou fonts occupy the private-use code points U+F17C1 to U+F17F6. In the standalone fonts, they are mapped into the ASCII range for easier access. GPLv3 with font embedding exception.

Icemap A code generator for static maps. This allows easier maintenance of otherwise hardcoded tables of Unicode data built into FontAnvil and IDSgrep. GPLv3.

IDSgrep A program for searching databases of tree structures stored in “extended ideographic description sequence” (EIDS) format. This is intended for searching kanji databases according to spatial structure using a powerful query syntax inspired by regular expressions. The need for such queries came up during Tsukurimashou development, but the package may well be useful in other contexts as well. Includes code to interface to several different character dictionaries, including structural data extracted from Tsukruimashou. GPLv3.

Kleknev A coarse-grained profiler for build systems, such as the Tsukurimashou build. This is still rather experimental. It basically consists of a wrapper program that Make can invoke instead of /bin/sh, which will (when told to do so by a special environment variable) collect a bunch of performance information that can be analysed by a separate reporting program to provide some insight into where all the time goes during a long build. See my Web log article at <http://ansuz.sooke.bc.ca/entry/260>.

OCR Freely usable vector versions of the standard OCR A and OCR B fonts, based on the work of Tor Lillqvist, Richard B. Wales, and Norbert Schwarz. License terms vary, but the fonts should all be freely usable.

バージョン管理システムからビルド Building from a version-control checkout

For the very latest public development version of Tsukurimashou and its associated tools, you can check out the source code from version control. Since the version control system stores only the source code as such, not any files that can be automatically generated, some additional steps are necessary to build from a checkout, beyond those used to build from a distribution.

Building from a source control checkout requires some additional dependencies not needed, or not so strictly needed, when building from a distribution package:

- ④ recent GNU Autotools, including in particular Automake at least 1.13;
- ④ the BuDDy binary decision diagrams library;
- ④ lex (preferably flex);
- ④ several files of Unicode character database information; and
- ④ if using the built-in support for automatically downloading the Unicode database files, then wget and unzip.

Check out the project by anonymous Subversion from OSDN Japan:

```
svn checkout https://svn.osdn.jp/svnroot/tsukurimashou/trunk \  
./tsukurimashou
```

Although there may still be a mirror of some version of Tsukurimashou on Github, it is out of date and unlikely to be updated. Use the Subversion repository.

The checkout command downloads the sources from the repository and creates a subdirectory called “tsukurimashou” in the current directory.

Use autoreconf to build the build system itself:

```
cd tsukurimashou
autoreconf -i
```

The `-i` option tells autoreconf to create missing files. This command is likely to produce many error and warning messages, but should work nonetheless if the Autotools versions are suitable. In particular, note that when autoreconf runs it may generate repeated “underquoted definition” warnings in relation to files in `/usr/share/aclocal` or similar. **This is the fault of some other package on your system;** it is not related to Tsukurimashou and should be harmless to the Tsukurimashou build. The issue is that many packages (IMLIB is one common culprit) install M4 files globally to provide customized Autoconf tests, and then if those files are buggy, autoreconf complains whenever it is invoked even if, as here, the buggy files are not actually being used.

Building Tsukurimashou requires FontAnvil (or FontForge); if you do not already have that installed, then the Tsukurimashou configure script will automatically build FontAnvil. However, some C source files for FontAnvil are not actually written under human control, but by the Icemap utility. If you were building from a distribution tarball, those files would be included ready-made; in a version control checkout, they must be built using Icemap, and that means Icemap in turn must be built. So the configure script will automatically build Icemap too, if you have neither FontAnvil nor Icemap already.

Running configure with no options should be able to generate a usable configuration if the necessary prerequisites are available, but it is likely to give some warnings suggesting better options.

```
./configure
```

Read any warnings produced, run `./configure --help` for more information or check other sections of this document, and if you wish, run the configure script again with your choice of other options. For instance:

```
./configure --enable-ot-features=all
```

The C sources built by Icemap for FontAnvil are based on Unicode character database information which must be downloaded from the Unicode Consortium; so there is an extra “make download” step (or manual equivalent) necessary before the main build, when building from a source control checkout in the Icemap, FontAnvil, Tsukurimashou scenario. Thus:

```
make download  
make
```

When `make` runs in a completely new checkout, it will probably go through a couple of repetitive cycles of re-building the Makefiles. This repetition should be harmless.

Icemap is also used by IDSGrep, and the IDSGrep configure script, unlike FontAnvil’s, will not speculatively configure the package to use an Icemap binary that does not yet exist. That means building IDSGrep in a completely fresh version control checkout, on a system where Icemap does not yet exist, will not work without additional human intervention. IDSGrep will, however, find and use a built Icemap in a sibling directory even if Icemap has not been installed on the system globally.

So to build IDSGrep from a version-control checkout, build Icemap first (either by itself or by doing a top-level Tsukurimashou build), and then run the IDSGrep configure (either by itself, or by running the top-level

configure a second time with an appropriate `--enable-parasites` setting), to allow `IDSgrep` to build with the freshly-compiled `Icemap`. These extra steps are not necessary when building `IDSgrep` from a distribution tarball, because the `Icemap`-generated sources will be included in the distribution automatically.

頒布のパッケージを作って Building distribution packages

It is not expected that general users of the `Tsukurimashou` project will have any reason to build distribution packages, and I don't encourage attempts because of possible confusion with "official" packages I prepare myself. However, it's worth having some documentation of the process for my own future reference and for the benefit of forks.

This project uses `Autotools`, and it is preferred to use `Autotools`'s "make distcheck" workflow for creating distribution packages. However, the make distcheck workflow depends on doing a `VPATH` build (that is, with generated files in a "build" directory separate from the source files) and that does not combine well with the multi-step information flow of `Tsukurimashou`. Unlike the typical C program that has clearly delineated source and object files, `Tsukurimashou`'s source code is often used to generate files that are themselves much like "source" code and are processed further, through several steps. These are somewhere between source and object code and don't naturally fit into the hard division contemplated by `Autotools`'s implementation of `VPATH` builds.

There are also many files, especially those related to `Icemap`-managed data tables in `IDSgrep` and `FontAnvil`, which may be more like source or more like object files depending on whether the local configuration is using `Icemap` or just using pre-made data tables obtained from a distribution package. To work well with `VPATH`, the system must adapt to those files being in different places depending on the current configuration. Different branches of the larger project (such as `FontAnvil` and `Icemap`) depending on each other, creates further complexity as the directory layout and build order for sub-packages change depending on the style of build.

The overall picture is that you shouldn't attempt a VPATH build except in the specific context of building a distribution package as described here. VPATH builds will cause warnings throughout the build system and it is assumed that if you ignore those warnings, then you think you know what you're doing. This is also my own small blow against the annoyance of other systems that produce dire warnings, or even refuse to compile at all, in **non-VPATH** configurations, because somebody had the idea that use of VPATH ought to be encouraged as a general practice for all builds.

Running `make distcheck` will run the test suites in those parts of the system that have them, and will refuse to package if any tests fail. That is, in general, the intended behaviour: we don't want to be distributing packages that don't work. However, FontAnvil's test suite is a work in progress. There are longstanding issues that cause some tests to fail; we don't want to remove those tests because they're an important guide to things that still need to be fixed; but they are basically harmless and should not impair releasing the entire project. Numerical issues in FontAnvil also cause some generated fonts to fail the strict validation of `fontlint.pe`, and that can cause several other parts of the system to fail testing.

The Tsukurimashou build system in general is intended to work on a wide variety of Unix-like systems; that is part of the point of using Autotools. However, the distribution packaging workflow in particular is not as thoroughly tested on a wide range of systems as the ordinary build process, and it may not work well on systems much different from my Slackware Linux installation.

To build a distribution of the entire project: start by configuring for a complete build, with all the font styles, features, character coverage, and parasites enabled. Do a "make download" to be sure all the necessary Unicode data is in place. Then run the complete build, in the source tree (not VPATH!). This should create usable and up-to-date versions of all the files that will eventually be distributed.

To run "make distcheck" for the entire project and have it work, it is necessary that the invocation of `configure` **inside** distcheck's build

process be one that will allow distcheck's build process to run and distcheck's invocation of the test suite to pass. That requires disabling the tests we expect to fail. Use the following command line.

```
make DISTCHECK_CONFIGURE_FLAGS='FONTLINT_FLAGS="-w 2,3,5,23"  
distcheck_hack=0.5pre distcheck
```

Cutting and pasting from the PDF file will probably not work because of the use of typographic quotation marks in this manual; you must edit them into ASCII quotation marks compatible with your shell.

The DISTCHECK_CONFIGURE_FLAGS variable passed to make specifies a command line fragment that distcheck will use when invoking configure. The FONTLINT_FLAGS variable (inside DISTCHECK_CONFIGURE_FLAGS) specifies error codes associated with FontAnvil's numerical problems, changing them into warnings instead of errors, so that they will not cause test suite failures.

There are other problems beyond fontlint errors in FontAnvil's own test suite, overridden by the distcheck_hack variable (outside DISTCHECK_CONFIGURE_FLAGS). This variable must be set to the current version number of the enclosed FontAnvil sources in order to take effect, so that the overrides cannot be accidentally applied to other releases without a deliberate intention to do so.

A standalone distribution of one of the parasite packages, that is, a package that can be used without the rest of Tsukurimashou, will usually need to contain some files that would be automatically generated (or at least copied) when building the whole project. For instance, the basic Metapost infrastructure code for generating fonts needs to be included in every package, but in a source control checkout there should be only one copy of it, and parasite packages automatically copy the master version to themselves when it is updated. Getting this right in a standalone package is a bit tricky.

To build a standalone distribution of a parasite package, start by building a complete version inside the project source tree. Run configure in the

parasite's directory with `--enable-parasite=yes` (or without this option, because it should be the default), and `build`. Then run `configure` again, with `--enable-parasite=no` and without doing a “make clean.” After that it should be possible to do “make distcheck” and have it create a working standalone package. Use the test-overriding options described for the whole-project package above, as necessary to get `distcheck` to stop complaining.

Once a standalone package is built, it should be possible to use it, and generate further distribution packages from it, without further reference to the rest of the Tsukurimashou sources, though stuff like rebuilding the Icemap-managed data tables will obviously require the relevant software.

ツールの文書化 Tool documentation

There are a number of scripts included in the `tools/` subdirectory. Most of these are intended only to be invoked automatically by the build system, but those intended to be invoked by hand are described below.

- ④ `autodep`: automatically maintain the header inclusion lines in METAFONT files, so that each file will include the files that define macros it uses, and only those. Files intended for use with `autodep` should include a comment line consisting of the case-sensitive text “% AUTODEPS” and any automatically-maintained input lines following that, terminating with a blank line. When the tool is invoked with the command line “`tools/autodep mp/ mp/*.mp`” from the root of the source tree, it will scan all the METAFONT-language sources, find definitions of macros, and where necessary rewrite the blocks of includes tagged with the special comment line in order to ensure that every file includes the definitions of all needed macros. There are a number of exceptions and special cases included to make it do the right thing with macros defined in multiple places as a result of style overrides, and so on.
- ④ `autolr`: automatically generate framework code for kanji that are

left-right combinations of existing components. Note that the code generated by this utility will in general **not** be usable as is; it is intended only as a starting point for human editing, to automate the repetitive work of getting character codes properly entered in decimal and hex, and so on. Run it with four arguments, from the root of the Tsukurimashou source tree: the (single character) component to put on the left; the prefix of the IDSgrep EIDS for dictionary entries with that component on the left; the (decimal numeric) X-coordinate for the division between the left and right sides; and the (decimal numeric) amount of overlap. The utility automatically runs IDSgrep to search for kanji with the specified structure and not already in Tsukurimashou, and generates to standard output some Metapost code that can be usefully edited to describe the new characters.

- ④ **build-plot**: generate a plot (using GNUplot, which is required) of the current progress of a full build. This was hacked in at the last minute to fill a specific need during the rather hurried packaging of version 0.9, and it may or may not be useful in any other contexts. Run it from the root build directory while a full build is in progress. It will create a file named `buildplot.png` which shows the current state of build progress. Computing the overall progress status of a Tsukurimashou build is complicated because even knowing which files need to be updated requires thousands of lines of code spread across several languages and systems; so this is an approximation. It reads the `makefile-vars.inc` file and attempts to identify all the Postscript, SFD, and OTF font files to be built in the current configuration. Then it checks the time stamps of all those files (if they exist at all) and computes what fraction of them have been updated since the current build began. It knows when the current build began by checking a tag file that is updated in the “kill-empty” step already wedged into an Autotools hook to execute at the start of a full build; thus, this will only work when the current build actually is a full build, initiated by just typing `make` rather than asking to build individual files. The completion estimate will also only be correct when all those files actually need to be built, that is, building from a fresh copy of the distribution or after `make clean`. It

will probably produce incorrect results if a build is not actually in progress (for instance, after a build has completed or been aborted). And the ETA estimate will usually be wrong. Despite all these limitations, it's a fun thing to run in a loop in a window to get some idea of where the build stands beyond the endless columns of MT1/PERL/RMO messages put out by `make`.

- `udcopyright`: check that all years in which each file was modified are included in its copyright notice. This assumes one has an SVN checkout in the current directory—intended to be a checkout from my private SVN repository, but it would probably work with one from the OSDN repository as well. It scans all SVN-controlled files for anything that looks like the project's standard copyright notice, and then reports on any files for which the set of years mentioned in the notice does not match the set of years in which SVN has log entries for that file. The idea is that once a modified version of the file has been checked into SVN in a given year, it has been “published” in that year and the year should be mentioned in the copyright notice; but when a new year comes we don't want to just go through and modify all files solely to insert another year in the copyright notice. By running this every so often, and manually updating the notices when appropriate, the hundreds of copyright notices in the project should painlessly converge on correctly reflecting the years in which files were actually changed.

For completeness, here is a summary of the functions of all the other scripts in the `tools/` subdirectory. Most of them are written in Perl.

- `add-flog`: attach the “font log” file to a FontAnvil SFD-format font; this requires re-encoding it to UTF-7 format.
- `livecode`: chase file inclusions and determine all the “live” (that is, actually invoked) macro definitions in a list of METAFONT-language files. The output of this is not (because lines may appear in the wrong order) actually valid code one could run; but the output of `livecode` is intended to have the property

that its checksum will almost always change when changes in the input necessitate recompilation, and will seldom change otherwise. Therefore the build system can use the checksum of `livecode`'s output for fine-grained change detection, to skip recompiling some files that (under GNU Make's timestamp-based scheme) would otherwise appear to require recompilation. As a special feature, intended for maintainer use only, `livecode` also supports a "find dead code" mode. If its first command-line argument is `"-d"`, then instead of its usual output, it will print a list of names of macros it thinks are **not** invoked. Running the dead code check against the entire code base may be a way to find superfluous METATYPE1 infrastructure that could be removed.

- ④ `make-ass`: generate the `assemble-font.pe` script, which in turn does most of the work of assembling 256-glyph subfonts to create full-coverage OTFs.
- ④ `make-book`: make decisions on how to split the multi-thousand-page proof document into smaller "books." The output is a set of `TEX` files which then get processed further to generate the actual PDFs.
- ④ `make-cfghl`: translate `configure` command-line strings into Hamlog for inclusion in `config.hl`, which then feeds back to control the build system.
- ④ `make-cover`: generate colour covers for the proof books.
- ④ `make-eids`: scan the proof files and generate an EIDS-format dictionary of character decompositions.
- ④ `make-fea`: perform text substitutions on "feature source code" (`fsc`) files to generate Adobe-format "feature" (`fea`) files expressing code-like font metadata.
- ④ `make-flog`: collect a bunch of debugging information and generate a font log file (see `add-flog` above).
- ④ `make-gpos`: compute information for GPOS features "mark" and "mkmk" by scanning the proof files (to determine anchor loca-

- tions and compatibility) and some generated FontAnvil scripts (to determine bearing adjustments made during kerning).
- ④ `make-hglpages`: generate METAFONT source code for the hangul precomposed syllable subfonts, which are all identical except for their starting indices and generated from a template.
 - ④ `make-kchart`: generate T_EX code for the kanji coverage chart.
 - ④ `make-kddata`: extract KANJIDIC2 data in Hamlog form.
 - ④ `make-name`: generate a font's NAME metadata table.
 - ④ `make-proof`: generate T_EX code for proofs and pretty-printed source code.
 - ④ `mp2pf`: stripped-down translation from Awk to Perl of a similar program in METATYPE1; this invokes Metapost in the appropriate way to generate something that can be processed into a Postscript font, and then does process the result into a Postscript font. Note that hinting support has been **removed** from this code, and the resulting Postscript fonts may not really be good or entirely format-valid; they are only intended as a way of getting the outline data into FontAnvil for later stages of processing.
 - ④ `mpdep`: generates a list of the dependencies of a METAFONT-language file (or a set of them), for automated Makefile generation. Compare with `autodep`, which actually edits the files to include the other files they should based on the macros they use; `mpdep` only looks at what other files each file does include, and reports that.
 - ④ `progress`: invoked during build to give progress reports and completion-time estimates.

ビルドのシステムの図示 Build system diagrams

<code>foo.bar</code>	foo.bar comes with the distribution
<code>naninani.dat</code>	naninani.dat is generated during build
<code>A → B</code>	A is input for B
<code>C → D</code>	C creates D
<code>E - - - - -> F</code>	E sometimes creates F

Figure 1: Legend for build system diagrams

This section is still under construction, but it is planned to be a graphical summary of what is going on in Tsukurimashou’s rather complicated Makefile. Figure 1 describes the symbols used, and Figure 2 shows the process by which fonts are built. Proportional fonts go through additional processing to generate the kern tables and so on, as shown in Figure 3.

Be aware that these diagrams are meant to clarify, not to formally describe, the code; this is **not** a formal modelling language like UML, and the semantics of the symbols are not necessarily entirely consistent. Also left out of the diagrams are language interpreters like Perl, and a great deal of information flow to and from the main Makefile itself. It passes configuration information, mostly in command-line arguments, to pretty much all parts of the system.

ハムログ Hamlog

Evaluating the consequences of build options like “--enable-chars” requires doing a certain amount of logical inference for which shell scripts are not well-suited. It might be possible to get GNU Make to do the necessary computations, inasmuch as it’s quite programmable, already required to build Tsukurimashou, and fundamentally a logical inference engine at heart. But that would probably involve creating many tiny files corresponding to logical propositions, which would waste space and cause other problems on some filesystems. A more elegant approach would be to use a real logic programming system, i.e. Prolog—which happens to be one of my research interests. But I didn’t want to create a

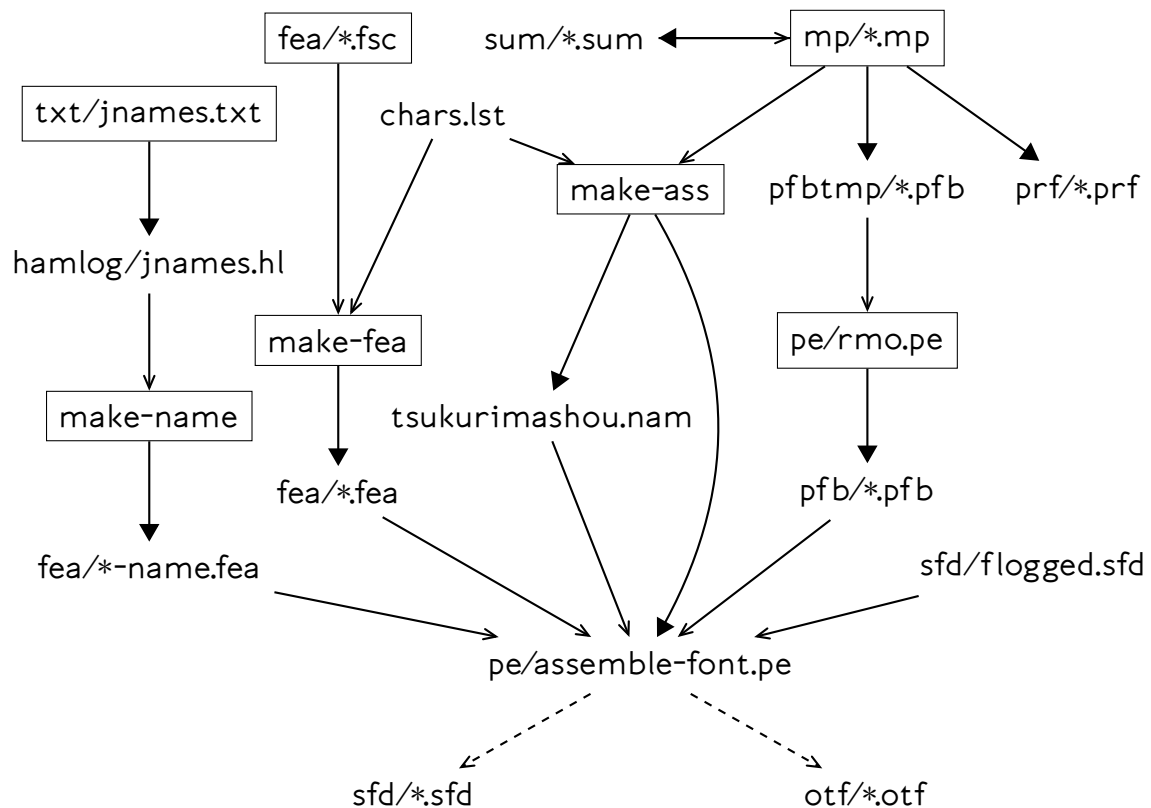


Figure 2: Building the fonts themselves. Checksums in `sum/` provide fine-grained change tracking on the Metapost files in `mp/`, which define Postscript “page” fonts each covering up to 256 code points. Those are processed to remove overlaps (RMO) and end up in `pfb/`. The `make-ass` script also scans the Metapost files to get a list of defined code points and some other metadata; its main output is a FontAnvil script that assembles the page fonts into full-coverage OpenType or FontAnvil native fonts. Non-proportional fonts become OpenType and are finished here; proportional fonts become FontAnvil native and are processed further. The `assemble-font` script also makes use of per-family feature files from `make-fea` (defining things like glyph composition) and per-font feature files from `make-name` (defining metadata table contents). The “proof” files in `prf/` are generated as a side effect of compiling the PostScript page fonts, and pass information about glyph construction to both the automatic documentation generator and the GPOS table generator.

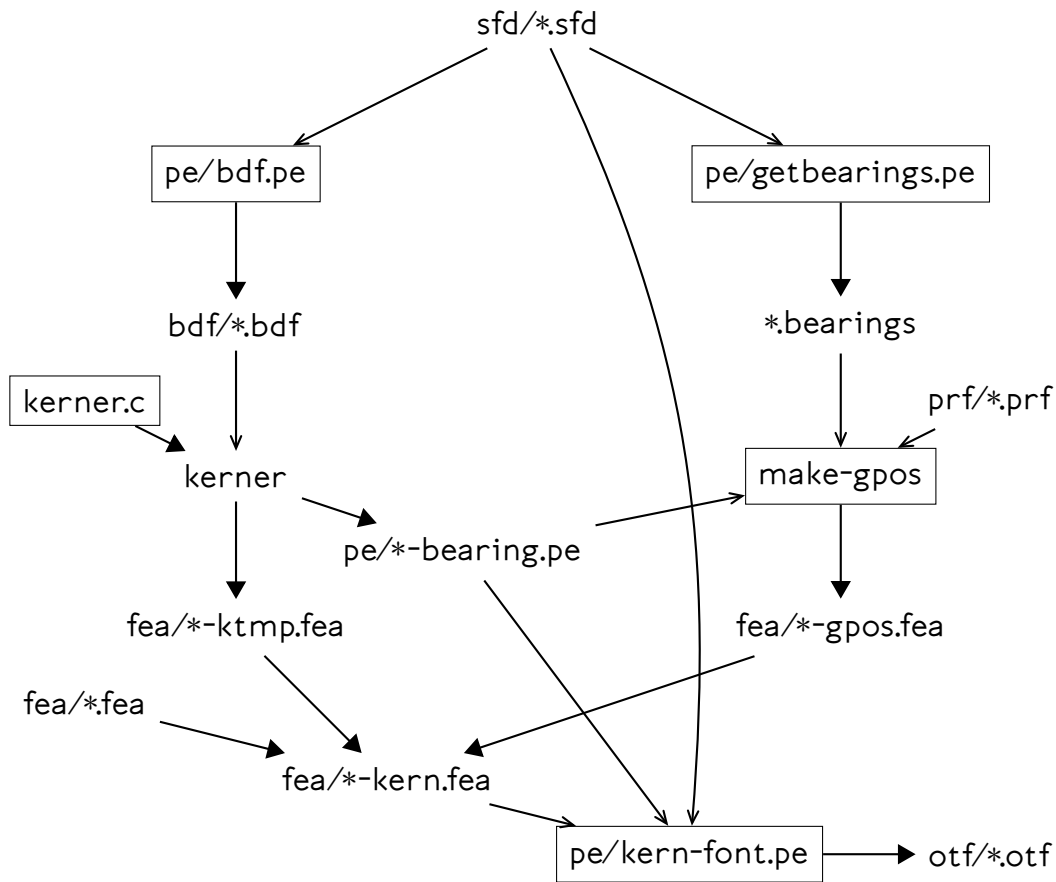


Figure 3: Additional processing for proportional fonts. The SFD file containing the font outlines is converted to a BDF bitmap font, which feeds the kerner program. That generates a “ktmp” feature file with the actual kerning data, and a FontAnvil script to apply the chosen bearings. At right, the make-gpos script reads the proof files to determine accent anchor locations. But since those must be adjusted when the bearings are set, it also needs “before” bearing information extracted by getbearings, and “after” obtained by reading the bearings script from kerner. The feature files for kerning, mark composition (GPOS), and family features are all combined and applied by the kern-font script, which runs the bearings script as a subroutine, finally producing a complete OTF font.

dependency on a Prolog interpreter because I think users will object to that; the existing dependencies of this package are already hard enough to sell. I also didn't want to bundle a Prolog interpreter, even though good ones are available on permissive licensing terms, because of the file size and build-system complexity consequences of bringing a bunch of compiled-language software into the Tsukurimashou distribution.

The solution: Tsukurimasho's build system will look for Prolog, and use it if possible. But the package also ships with something called Hamlog, which is a toy Prolog-like logic programming system written in Perl. (A ham is like a pro, but less so.) If the build system can't find a Prolog interpreter, it will use Hamlog instead. Hamlog is slow, and internally horrifying, but it works in this particular application. It is not particularly recommended for any other applications.

The configure script looks for SWI-Prolog, ECLⁱPS^e-CLP, and GNU Prolog. If one of these can be safely detected, it will be used. ECLⁱPS^e-CLP has the issue that it shares a name with a widely-used programmer's IDE, so it is not safe for the configure script to actually execute an executable called "eclipse" if it finds one. If something that might be ECLⁱPS^e-CLP is detected, the configure script puts up a warning and the user is free to enable it explicitly.

The rest of this subsection can and probably should be skipped by anyone who isn't both a Perl and a Prolog hacker.

Still with me? The way Hamlog works is sort of fun and so I'm going to spend a few pages describing it for those who are interested, if only as an example of something you should Not Try At Home. The idea is to use regular expressions for the operation of finding a clause whose head matches the current goal. Hamlog reads its program into a Perl hash, where the key is the functor and arity (like "foo/3") and the value is a newline-separated pile of clauses in more or less plain text. When it tries to satisfy a goal, it takes the goal (which starts out as plain text) and converts it to a regular expression that will match any appropriate lines in the database. Variables in the goal turn into wildcard regular expressions; ground terms turn into literal text; and then when there's a match, parenthesized sub-expressions are used to

extract the body of that clause.

It is because of the use of regular expressions that Hamlog doesn't do compound terms, and in turn is likely not Turing-complete (though I haven't thought carefully through all the possibilities of using recursive predicates or atom splitting to build stack-like data structures). As all the world knows, it is impossible to write a regular expression to match balanced parentheses. Current versions of Perl actually bend the theory with experimental extensions that do allow the matching of balanced parentheses, so that in a certain important sense Perl regular expressions **are not regular expressions anymore at all**, but even I am not quite twisted enough to actually deploy such code. Things in Hamlog that look like compound terms (such as the sub-goals in a clause body) are handled as special cases; but the point is that arguments to a functor that will be used as a goal have to be either atoms or variables. This also means Hamlog doesn't do Prolog syntactic sugar things that expand to compound terms, such as square brackets for lists.

Once there's a match, it does string substitution on the matching head, the current partially-completed goal, and the body, to get a new modified body for the clause, taking into account any variables assigned by the head match. The new clause body gets substituted into the current partially-completed goal (which is a string) as a replacement for the head that just matched. So the partially-completed goal is a sort of stack of comma-separated heads that grows from right to left and implicitly contains all the assigned variables.

Because of the simplistic way variables are given their values, it is dangerous to use the same variable more than once in the same head, so constructions like "foo(X,Y,X)" should not be used. If you want to do that you should instead write "foo(X,Y,Z):- =(X,Z)." Note the non-sugary use of "=" as a functor, since the more common infix notation isn't supported. Note also that there should be a space between ":" and "="; Hamlog doesn't require that but it may reduce the likelihood of parsing problems should the same code be run on interpreters other than Hamlog.

Variables in clause bodies are renamed once (using the clause serial

number), when the clauses are loaded; as a result if the same clause body gets expanded a second time while variables from its earlier expansion are still unassigned, there could be trouble. This is not a very likely scenario, but it's worth mentioning.

Clauses in the database have serial numbers; and when a choice point goes on the stack, the serial number of the clause at which it matched is part of the record on the stack. Then if the interpreter backtracks to re-satisfy a clause, it writes the regular expression in such a way that it can match all and only serial numbers greater than the last place it matched. Creating an "integer greater than N" regular expression was surprisingly difficult—it's a simple enough concept but there are several cases that all must be handled properly or weird bugs turn up.

Syntax is simplified from Prolog. Variables start with an uppercase letter or an underscore and may contain uppercase alphanumerics and underscores. Atoms start with a lowercase letter or numeric digit and may contain lowercase alphanumerics and underscores. For Prolog compatibility, atoms starting with digits should not contain anything other than digits, and the only atom starting with zero that should be used is zero itself; but Hamlog doesn't care about those points. Things containing a mixture of upper- and lowercase alphabetic characters should not be used. The special tokens "!" and "=" are technically treated as atoms too, but you should only use them in their typical meanings of cut and unification, and "=" should only be used with the general prefix syntax applicable to all functors, not as an infix operator (see above).

Variable names starting with, and especially the unique variable name consisting entirely of, an underscore, are not special in Hamlog. Beware, that means "foo(,_)" contains only one variable occurring twice when interpreted as Hamlog, not two distinct variables as in Prolog, and it violates the "only one appearance of each variable in a head" guideline. The unique variable name consisting of one underscore is probably best avoided entirely. But it may be desirable to use variable names starting with underscores anyway in some cases, because of their specialness to Prolog interpreters. I was really tempted to allow and use arbitrary UTF-8 (in particular, kanji) in atoms but refrained because of the desire

for Hamlog code to be easily readable by nearly all Prolog interpreters.

I tried to keep the number of built-in predicates to an absolute minimum, partly because any that are not standard Prolog have to be re-implemented in Prolog (and probably once for every supported Prolog) to build the shell that executes Hamlog programs on a Prolog interpreter. Here's an exhaustive list.

- ☛ ! [cut]. This is implemented by string substitution as well: when a clause body gets added to the to-satisfy stack, there's an additional regular expression substitution pass that converts any instances of !/0 in the body into !/1 where the argument is an integer identifying the current height of the choice-point stack. If at any point in the future we attempt to satisfy a !/1 goal, then the stack gets popped back to that point (discarding any choice points created between the time the ! got its argument and the present time). For this reason, ! should not be used as an atom or functor for any other purposes than as the cut, even if to do so would otherwise be valid Prolog.
- ☛ fail, which causes immediate backtrack (useful in conjunction with cut to implement negation).
- ☛ true, which is not actually used, so maybe I should delete it.
- ☛ var, true if the argument is a variable (not yet bound to an atom). This is important in Hamlog because many predicates need to be written to accept more than one instantiation pattern for their arguments.
- ☛ atom, true if the argument is an atom. This is implemented by rewriting the database entry for atom/1 on the fly; when you call "atom(foo)" it magically changes to having been defined as either the single fact "atom(foo)." or nothing, depending on whether "foo" is an atom.
- ☛ atom_final(AZ,A,Z), where AZ is an atom whose last character is Z and A is everything except the last character. Used for building

and tearing apart atoms like “page00.” This requires some careful handling in other interpreters because Hamlog has no concept of quotation marks and treats single-digit integers exactly the same as atoms whose names are the ASCII digits; real Prologs have more subtle type handling. As with `atom/1`, this is implemented by magically creating appropriate clauses on the fly.

- ☹ `=/2`. This also creates clauses on the fly. At least one of the two arguments should already be atomic when the goal executes, though that is rarely difficult to guarantee in practice.

And that’s it for built-in predicates. Note that goals in a clause body also can only be combined with comma for conjunction (no semicolons for disjunction, and without them parentheses become unnecessary and are not supported either). There is also no syntactic support for negation. However, you can (and the existing code does) compute negation and disjunction using multi-clause predicates, `cut`, and `fail`. What you can’t build is any kind of I/O—so how can Hamlog programs communicate with the outside world?

The interpreter (after checking for the “`--version`” and “`--debug`” options, which do fairly obvious things) interprets its first two command-line arguments as a template and a query. The template ought to be a valid Prolog compound term for compatibility with other interpreters, but Hamlog actually treats it as a string. Then it backtracks through all solutions to the query, attempting to instantiate all variables, and writes (newline separated to standard out) all the **distinct** values assumed by the template. This is basically the same operation as Prolog `findall/3` followed by `sort/2`, which is how the Prolog shell for Hamlog implements it. Any remaining command-line arguments, and standard input if there are none of those, will be read in the usual `<>` way to fill the Hamlog program database. Hamlog code is conventionally stored in “*.hl” files.

In the build system, the string of comma-separated tags for things like characters to be selected gets converted (by the “`make-cfghl`” Perl script) into a few clauses of Hamlog and written into the file `config.hl`. Also written there is a list of `page_exists/1` facts naming the 256-code-point

blocks for source files that exist in the mp directory. Then elsewhere in the build system, it invokes Hamlog with appropriate queries against config.hl and select.hl to get lists of characters, OpenType features, and other things that the user does or doesn't want, based on knowledge built into select.hl of what the different selection tags actually mean.

It is planned that in a future version, the KANJIDIC2 file will be automatically translated to more Hamlog facts expressing which kanji are or aren't included in the different grade levels; then it will be possible to use options like "--enable-chars=kanji,-grade3' for finer-grained selection of kanji.

In the case of a Prolog interpreter other than Hamlog, there is some other code written in that interpreter's own language to allow it to execute Hamlog programs and export something resembling this command-line interface to the Makefiles. Since Hamlog programs are also syntactically valid Prolog, this support shouldn't be difficult in general. See the swi-ham.pl file for what currently exists of this nature. The main advantage of using a non-Hamlog interpreter is simply speed.

カーニングしかた

Kerning

This is a summary of how the automated proportional spacing and kerning code works.

First, the build system generates the future PS font as an OpenType file, all complete except for widths and kerning. It then calls `bdf.pe`, which sets all the bearings to 50 font units and makes a BDF-format bitmap font scaled so that the reference kanji square (1000 font units) takes up 100 pixels. This script filters out glyphs that should not be subject to kerning, which includes those with zero width (mostly combining characters), all the hangul individual-jamo code points (whose layout is handled by OpenType substitution features, and kerning would just make it too complicated), and the special glyphs used by vertical fraction composition (much the same situation as the hangul jamo).

The “`kerner`” C program reads that BDF font, puts all the glyphs into a common bounding box big enough to contain any of them, and finds the left and right contours—basically, the x-coordinates of the leftmost and rightmost black pixels on each row—for each glyph, as well as the margins, which are defined as the x-coordinates of the leftmost and rightmost pixels on **any** row of the glyph.

There is some special processing applied to the contours to make them more suitable. Consider what happens in a glyph like “=”: many horizontal rows, namely those above and below the entire glyph and those in between the two lines, contain no black pixels at all, and so the leftmost and rightmost black pixels in those rows are formally undefined. If we set it next to another glyph like “.” which only has ink in rows where “=” does not, then a kerning algorithm that looked only horizontally might let the period slide all the way under the equals and out on the other side. There has to be some vertical effect to prevent that. So the Tsukurimashou kerning program makes a couple of passes over each contour (one forward, one backward) to enforce the following rules, which (except for glyphs containing no black pixels at all, which are removed from consideration) fully define where the

contour should be.

- ④ The right contour cannot be any further left than the rightmost black pixel in the row.
- ④ The right contour cannot be more than 10 font units (one pixel) further left than its value in the next or the previous row.
- ④ If the right contour in the next or previous row is left of the **left** margin, then the right contour in this row cannot be more than 3 font units further left than in the next or previous row.
- ④ Subject to the above rules, the right contour is as far left as possible.
- ④ The left contour's rules are the mirror image of these.

These rules can be imagined as simulating something like the way letterpress printers kern type by physically cutting the metal type bodies at an angle to fit them more tightly together: it's bad to cut off part of the actual printing surface; you basically cut at a 45° angle; but with glyphs that only have a small vertical extent, so that the 45° angle would cut all the way across to the other side, then you want to use a more vertical angle (in this case, 16.7° from vertical) so you don't end up setting the next character actually earlier on the line. Figure 4 shows a typical example of the contour computation.

The right margin is subtracted from the right contours and the left margin from the left contours to get, for each glyph, a vector of numbers describing its shape along the left and right sides, independent of the width of the glyph.

All the left contour vectors, and (independently) all the right contour vectors, are subjected to a modified k-means classification. Initially, the contours are put in 200 classes, according to the values of a simple hash function applied to the glyph names. This is a change from earlier releases in which a round-robin was used: the advantage of the hash




Figure 4: Glyph contour computation

function is that although it remains deterministic, it helps break up a phenomenon that tended to happen with the Genjimon glyphs, where because of excessive symmetry in the initial arrangement, the classifier could never put the similar glyphs together.

Then for each contour the program asks the question “How far is this contour from the centroid of its class, and if I moved it to a different class, how far (after accounting for the fact that the centroid changes when I add the glyph) would it be from the centroid of the new class?” If moving the glyph to some other class would make it closer to the new centroid, then the glyph gets moved to the other class where its distance to the centroid will be minimized. Note that a glyph in a class by itself will never want to move out of that class, because its distance to the centroid is already zero.

There is an extra rule that the classification will never move a glyph into a class in such a way as to make a class larger than the “class limit,” defined to be the larger of 100 glyphs, and three times the total number of glyphs divided by the total number of classes. The purpose

of this rule is to counteract a tendency seen in some experiments for the classifier to create a few huge classes (for instance, a single class containing a large fraction of the 11172 precomposed Korean syllables) that cause font subtable size problems. It is not clear just where the limits are on how big a class may be, but this limit appears to work at the moment.

Glyphs are examined in this way until no more such moves are possible. The idea is that at the end of it, the glyphs will all be in classes that are as tightly clustered as possible. It's not guaranteed to be a global optimum (in other words, it's possible that some other assignment of glyphs to classes might be better; really optimizing this problem is difficult) but it's guaranteed to be a local optimum in the sense that it can't be improved by changing the assignment of just one glyph, and it's expected to be pretty good overall. Note that the initial assignment was deterministic (where random would be more usual for this kind of algorithm) because it seems undesirable for the kerning distances, overall, to be non-deterministic; my copy of the font shouldn't have different metrics from yours if they were compiled from the same sources with the same options.

After classification, we've got 200 classes of left contours and 200 classes of right contours. The actual kerning is done class-to-class, using the class centroids, rather than glyph-to-glyph. That way we will end up with up to 40000 kern pairs instead of millions. OpenType supports this kind of kerning pretty well. There will be a feature file generated listing the contents of the classes and the distances for each pair of classes, and that's much more efficient both in source and compiled form than specifying a distance for every pair of glyphs.

The number 200 (up from 150 in earlier versions) was chosen by educated trial and error. Every glyph must be in a class for full kerning; but no class can be too big or FontAnvil barfs. That in turn limits the size of the average class to some maximum, and so limits the number of classes to some minimum. Using 200 classes seems to be enough for the current Mandeubside fonts, which are roughly 13000 glyphs each, and it is reasonable to estimate that none of the fonts created in the planned scope of this project will have much more than

that many glyphs.

To kern two contours together, we can compute a closeness value for each row by saying “if we positioned the margins of the glyphs this much apart, how far apart would the contours be in this row?” That distance, divided by a constant representing the optimal distance (currently 230 font units) and raised to a power representing how much extra weight to give to the closest points (currently 3), represents closeness for the row. The sum of closeness for all the rows would be equal to the number of rows in the case of two perfectly vertical lines 230 units apart. The kerner program adjusts the margin-to-margin distance so that the sum of closeness is equal to that. It uses a binary search to do that adjustment, which is probably not optimal for a fixed exponent (there should be an exact analytic solution possible without iterating) but has the big advantage of not requiring a redesign should the exponent or even the entire closeness function change.

The closeness computation is shown schematically in Figure 5. Note that this closeness is calculated on the contours, as defined previously, rather than the actual shapes of the glyphs; it is also done on the centroids of the kerning groups, thus generic contours each representing many glyphs, rather than the contours of any specific individual glyphs.

The effect of the exponent 3 in the calculation is to give much more weight to points that are close together, as suggested by the shading in the figure. If we’re kerning a pair like “]<”, we want to pay more attention to the point of the less-than than to the distant ends. An exponent of 3 means that points at half the distance count eight times as much toward overall closeness, so there’s a strong bias toward seeing the points of closest approach. If we imagined using a larger exponent, this bias would be even stronger; in the limit, with an infinite exponent, the kerning would be determined solely by setting the closest approach to the optimum without reference to any other points. That is how most auto-kerning software works; but the results tend not to be good because in a serif font with a pair like “AV,” inserting the ideal vertical-to-vertical distance between the serifs is going to place the stems, which are much more visually important, too far apart. Using an exponent somewhat less than infinity causes the stems to still have




Figure 5: Closeness computation

some significant weight. The value 3 was chosen by trial and error and may be subject to further adjustment.

Once all the class-to-class distances have been chosen, it remains to choose the bearings for the characters. Recall that kerning was computed from margin to margin, that is the amount of space to insert between the strict bounding boxes of the glyphs. Adding a certain amount of extra space to the glyphs themselves, and subtracting it from the kern distances, may result in a better, more concise kern table, as well as better results when glyphs from this font are set next to spaces, things from other fonts, and so on.

The first step is that the kerner program finds the average of all kern table values, and puts half that much bearing space on either side of every glyph, adjusting the kern values accordingly. This has the effect of giving every glyph an “average” amount of space, and changing the overall average kern adjustment to zero. If we were to throw away the kerning table at this point and just use the bearings, then in some sense these bearings would give the best possible approximation of the

discarded kerning table.

Then for each left-class (which is actually a class of **right** contours: it is a class of glyphs that can appear on the left in a kerning pair) the program finds the maximum, that is farthest apart, amount of kerning between that left-class and any right-class. Two thirds of that kerning amount are added to the right-bearing of the left-class. The concept here is that we generally want kerning to be pushing things together, not pulling them apart, so the “default” amount of kerning indicated by the bearing should be near the maximum distance apart, from which individual entries can then push things closer. Also, we generally want most (two thirds) of this adjustment to happen to the right bearing of the left-hand glyph in the pair.

Then to clean up the rest, the program examines each right-class (which is a class of left contours) and similarly finds the furthest-apart kern pair and adds that to the left bearing of the right class, adjusting all kern pairs appropriately. At this stage it's guaranteed that all the kern table entries will be zero or negative: kerning only pushes glyphs together from where they would otherwise be, it never pulls them apart.

Kern table entries are dropped if they are less than ten font units; that cuts the size of the table considerably. In a change from earlier versions, the kern values are not otherwise rounded (beyond being integers). The table is written to a fragment of an Adobe .fea file, with subtable breaks on left-class boundaries each time a subtable grows past 5000 entries; that means subtables actually end up a little over 5000 entries each. That seems to be how big they can get without overflowing the OTF table-size limits. Bearings are written to a PE script.

The build system runs kern-font.pe, which applies the output of the kerner program to the font. Something else that kern-font.pe does is to add a hardcoded additional bearing of 40 on the left and 80 on the right to all Japanese-script characters (kana and kanji); by trial and error, this seems to make the results look better. It seems to be simply a fact that Japanese characters need more space between

them to look right than Latin characters do at the same type size. Something similar should probably be done to Korean-script characters, but that has not been determined yet.

That is how the horizontal spacing of the font is currently computed. It still isn't perfect, but some progress has been made.

さてさてなにができるかな？