KanjiCompass: An Etymology-Driven Adaptive Kanji Learning Tool

Sigrid L. Klinger[a] and Sven Strickroth[b]
LMU Munich, Munich, Germany
{*sigrid.klinger, sven.strickroth*}*@ifi.lmu.de*

Keywords: Etymology-Based Learning, Self-Regulated Learning, Graph-Based Interfaces, Kanji Learning.

Abstract: Learning Kanji is a complex and critical component of Japanese language acquisition, requiring learners to understand its semantics, morphology, and phonology. Traditional rote memorization methods often overlook Kanji's etymological and structural nuances, limiting their effectiveness. This paper presents an etymology-driven, adaptive Kanji learning tool designed to visualize Kanji relationships, reduce cognitive load, and enhance learner engagement. The tool features interactive graph visualizations, personalized learning recommendations, and integration with Anki flashcards for explorative, self-regulated learning (SRL). The tool was evaluated for its usability and adaptivity in a field study with 19 participants. Overall, the tool's usability was well-received, with the detailed Kanji graph and Anki integration being commended for their clarity and ease of use. Personalized learning recommendations were particularly valued for providing adaptive and targeted learning paths. However, the macro-level perspective provided by the overall graph was found overwhelming by some users. Results also indicate that learning goal motivation strongly influenced engagement, with motivated users benefiting more from the tool's adaptive features. Key contributions include methods for visualizing interconnected knowledge, recommendations for personalized learning paths, and supporting tools for encoding and retrieval stages.

1 INTRODUCTION

In 2021, about 3.79 million people worldwide studied Japanese formally, plus countless self-learners (The Japan Foundation, 2023). A crucial and challenging component of mastering Japanese is its writing system, particularly the use of Kanji, which plays a fundamental role in both literacy and deeper language comprehension.

Kanji, a morphographic script integral to the Japanese writing system, poses significant challenges to learners (Rose, 2017). In contrast to a phonographic script like the Latin alphabet, where every character represents a sound, characters of a morphographic script represent a meaning. There are 2,136 Kanji for regular-use issued by the Japanese government, so-called jōyō Kanji, which are deemed essential for functional literacy in Japanese, covering the vast majority of Kanji seen in everyday texts. Of these, the 500 most frequent characters account for 75 % of all Kanji occurrences in written Japanese (Crowley, 1968). The script's complexity lies in its multilayered composition of semantics, morphology, and phonology, making it distinct from other writing systems. As a simple example the Kanji 土 has the meanings of 'soil, earth, ground, Turkey'[1] (semantics), looks similar to another Kanji 士 (morphology), and can be pronounced "tsuchi", "do", and "to" (phonology). Even though Kanji are rooted in Chinese characters, they differ significantly. One Kanji can have several pronunciations, several meanings, and is normally combined with other Kanji or Hiragana—another Japanese script—to form words. This combination of factors needs to be considered when learning Kanji, making it impossible to learn Kanji as one would learn vocabulary. Unlike vocabulary, where one usually only learns one meaning and pronunciation, Kanji often have several meanings and pronunciation options that depend on the context.

While some materials integrate structural or etymological explanations, many Kanji learning resources still emphasize rote memorization, which has been criticized for limiting the learner's ability to grasp the structural nuances of Kanji (Mori, 2014). Existing methods, such as visual imagery or narrative aids, attempt to enhance memorization but often overlook the potential of etymology as a pedagogical tool (Rose, 2017). By leveraging etymology, where the origin and development of a character is explored, learners can

[a] https://orcid.org/0009-0006-9277-8355
[b] https://orcid.org/0000-0002-9647-300X

[1] https://jisho.org/, last accessed: 2025-02-11

Klinger, S. L. and Strickroth, S.
KanjiCompass: An Etymology-Driven Adaptive Kanji Learning Tool.
DOI: 10.5220/0013341000003932
In *Proceedings of the 17th International Conference on Computer Supported Education (CSEDU 2025) - Volume 1*, pages 472-483
ISBN: 978-989-758-746-7; ISSN: 2184-5026

understand not only the meanings and pronunciations of individual Kanji but also discern meaningful relationships among characters, fostering a more cohesive learning experience.

Etymology not only enables learners to grasp the meanings and pronunciations of individual Kanji, but also facilitates the identification of connections between known and unknown characters. These inherent interrelations between Kanji, often disregarded in conventional teaching materials, are nonetheless recognized by learners, who frequently form their own assumptions about these links (Rose, 2017). Unfortunately, these assumptions can lead to misconceptions and frustration, diminishing motivation and the overall learning experience.

To address these challenges, this paper introduces an etymology-driven, adaptive learning tool, which answers the following research questions:

- How can the structure of a Kanji itself and its relation to other Kanji be visualized in a learning tool to minimize the cognitive load needed to learn Kanji?
- How can adaptive recommendations for learning paths be given to guide learners in deciding which Kanji to learn?
- How do learners perceive and use the tool's visualization and recommendation features, and to what extent do these features motivate and engage them in exploring and deeply learning Kanji?

By emphasizing the interconnectedness of Kanji through their etymological roots, the tool aims to provide both, self-learners and formal students, with a deeper understanding of Kanji, reducing cognitive load and preventing erroneous assumptions. This technological solution presents an opportunity to shift the focus from isolated memorization to a more integrated, explorative learning experience, better suited to the complex nature of Kanji acquisition. The tool is not intended to replace conventional and established learning methods but to complement them. The contributions are: (1) a purposeful layout with navigation elements aimed at minimizing cognitive load and optimizing the learning process, (2) an approach to visualizing the interconnectedness of a learning subject, and (3) a method for providing personalized learning paths by giving unobtrusive learning recommendations.

The paper is structured as follows: Section 2 provides background on Kanji learning, Section 3 reviews related work, Section 4 presents the prototype, Section 5 details evaluation and results, followed by discussion and future research.

2 BACKGROUND

To provide an understanding of the current state of Kanji learning this section will provide an overview on Kanji and research concerning Kanji learning methods. The Japanese writing system is comprised of four scripts: Hiragana, Katakana, Rōmaji and Kanji. Hiragana and Katakana each consist of 45 unique characters, while Rōmaji is used to transcribe Japanese words with the Latin alphabet. Kanji, however, do not have a set number of characters. Some online dictionaries comprise more than 28,400 Kanji[2], but for learners, the 2,136 jōyō Kanji are of primary importance. To be proficient in Kanji, one needs to know the meaning of a Kanji, its pronunciation, and the ability to recognize and write it.

Kanji can be classified into six etymological categories based on their historical origins, with three major categories alone accounting for 97.5 % (Tamaoka et al., 2017) of all Kanji. This etymological perspective is valuable because it not only aids memory by anchoring characters in meaningful contexts but also helps learners recognize patterns and connections among Kanji. Examples of the three main categories include: 高 ("tall"), originally depicting a tower (pictorial), 森 ("forest"), formed by three "tree" (木) components (metaphorical), 際 ("border," pronounced *sai*), combining 阝 ("hill") and 祭 (also *sai*) (phonological).

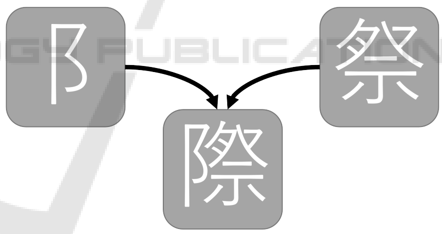

Figure 1: A visualization of the two graphemes 阝 and 祭 forming the Kanji 際.

While some characters represent a single entity, most are composed of multiple graphemes, including one radical. A grapheme is "the smallest unit in a system of writing a language that can express a difference in sound or meaning".[3] An example in the English writing system are the words "shake" and "snake" only differentiable by the graphemes "h" and "n". Radicals are foundational building blocks that help make Kanji searchable and provide semantic clues, while graphemes contribute directly to a character's meaning or phonetic reading. This leads to a natural parent-

[2]https://kanji.jitenon.jp/, last accessed: 2025-02-11

[3]https://dictionary.cambridge.org/dictionary/english/, last accessed: 2025-02-11

child relationship between graphemes and the Kanji they form, with pictorial Kanji and radicals as roots.

Despite Kanji's compositional structure, many learning materials like JiShop[4] break them into their smallest components without considering etymological significance. This poses a challenge for research and pedagogy: distinguishing between graphemes, which contribute to a Kanji's meaning or reading, and components, which may not. For instance, the Kanji 際 (pronounced "sai") contains the phonetic grapheme 祭 (also "sai") and the radical 阝, but can also be broken into smaller components (丿, 二, 小, 阝, 癶, 示)[5], which offer no phonetic and no mnemonic value. Distinguishing between etymologically meaningful graphemes and irrelevant components is therefore key to optimizing Kanji learning.

Some Kanji were simplified in the 1946 reform, resulting in "old" versions being replaced with standardized forms. For instance, 国 ("country") was historically written as 國, yet both versions belong to the same etymological category of metaphors: 囗 ("border") enclosing 或 ("estate") or 玉 ((king's)"jewel"). This demonstrates that an etymological analysis of modern Kanji remains useful for learners, even without referring explicitly to older forms. The structural changes often preserve underlying semantic and phonetic relationships, making etymology a valuable tool for Kanji learning.

For effective learning, it is important to differentiate the three stages of memory: encoding, storage, and retrieval. Understanding these stages is crucial, as they highlight different approaches to Kanji learning—from initial character recognition to long-term recall. Learning methods for the encoding stage include visual associations, components analysis, and mnemonic sentences (Rose, 2017). The retrieval stage involves methods such as rote memorization or repeated writing of characters. These methods are incorporated in some learning materials. "Kanji Look and Learn" (Banno et al., 2009), for example, focuses solely on visual aids, while Heisig (1994) relies entirely on mnemonic phrases. These materials often present their own learning method as "the one and only easy method" to learn all Kanji, yet they typically overlook the etymological creation of characters.

This overview highlights the complexity of Kanji and the limitations of traditional methods, emphasizing the need for an etymology-driven approach that can make learning both more efficient and meaningful.

[4] http://www.jishop.com/, last accessed: 2025-02-11

[5] https://jisho.org/, last accessed: 2025-02-11

3 RELATED WORK

This chapter explores Kanji learning tools, beginning with graph-based representations to illustrate character relationships, followed by flashcard systems for memorization, and concluding with adaptive learning technologies for personalized support. Technologies for language learning encompass a wide range of approaches, but most focus on vocabulary or grammar (Cook, 2016) and are not readily applicable to Kanji learning. Kanji learning requires tools that address its unique characteristics, such as etymological structure and compositional relationships.

As described, Kanji are either composed of graphemes or are standalone, indivisible characters. The structure of Kanji follows a hierarchy, with standalone Kanji forming the top level, followed by combinations of standalone and/or composite Kanji. This hierarchy can be represented by a directed graph, allowing learners to better understand the relationships between Kanji (Komarek et al., 2015). In this context, the question arises of how to build and visualize a Kanji network to support learning. Existing research on graph models of Chinese (Hànzì) and Japanese characters offers initial approaches, as Kanji and Hànzì share structural similarities.

Li and Zhou (2007) modeled a Hànzì network with 6,652 characters and 1,624 components as nodes, analyzing its density, path length, clustering, degree distribution, and assortativity. They found that the network exhibits complex small-world properties. However, a visualization of the entire network was found too complex to be legible. Jeronimus et al. (2017) examined a network of 1,945 Kanji, showing that some Kanji had very high degrees due to extensive breakdown into components and not graphemes. They also visualized only a partial graph.

Yan et al. (2013) developed a Hànzì network to create a learning sequence, finding that a few characters construct many others, while most characters form only a few. This network utilizes color to indicate compositional features and node size to suggest learning order based on stroke count (cf. Figure 2). However, while this approach considers the graphemes of the characters, it does not account for how learners can internalize these relationships, nor does it adapt to the learner's existing knowledge base. The compositional relationships are presented statically, and further explanation is necessary for learners to fully understand the characters' structure and meaning. Although Yan et al. (2013) addresses some aspects of character composition, their network lacks the adaptability necessary to support personalized learning pathways. Kovacs[6]

[6] https://thekanjimap.com/, last accessed: 2025-02-11

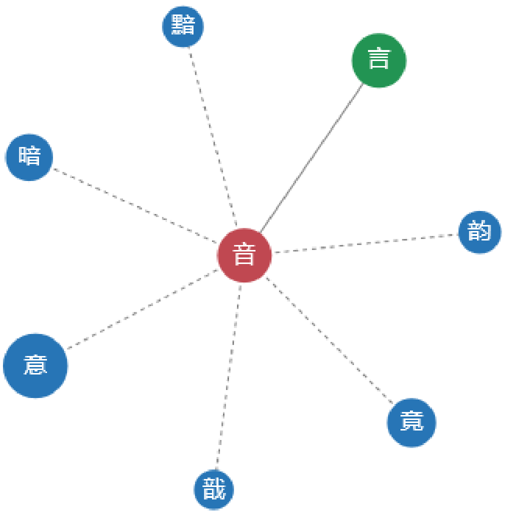

Figure 2: Network graph of the Hànzì 音 from the “Meaningfully learning Chinese characters” tool by Yan et al. (2013).

created the web application *The Kanji Map*, using a force-directed layout to visualize connections for a given Kanji. However, this visualization does not account for the etymological structure of Kanji, limiting its effectiveness for learning.

In language learning, flashcard programs have been widely recognized as effective tools for memorization (Hanson and Brown, 2020). A widely used flashcard program, continuously updated since 2008, is Anki.[7] Anki (暗記, “memorization”) is an open-source flashcard software originally developed by a Japanese language learner for studying Japanese and English but became a generic tool for learning with flashcards. Anki still remains popular in the Japanese learning community due to its many add-ons, tutorials, and active user base. Users can create custom multimedia flashcards with a front and back, customizable via HTML and CSS. Anki’s algorithm schedules review sessions using spaced repetition, based on the SuperMemo 2 (Wozniak, 1990) and the extended version *Free Spaced Repetition Scheduler* algorithms, though users can adjust intervals manually.[8] When a card transitions from *learning* to *review* status, it is assigned an *ease factor*, which reflects the learning difficulty level. This factor starts at 250 % and can drop to a minimum of 130 %, depending on past responses and repetition frequency. Zimmerman and McMeekin (2020) suggest that educators should recognize learners’ preference for simple flashcard programs when learning Kanji and, thus, avoid investing significant effort in creating learning games.

[7]https://apps.ankiweb.net/, last accessed: 2025-02-11

[8]https://faqs.ankiweb.net/what-spaced-repetition-algorithm.html, last accessed: 2025-02-11

One approach to implementing personalized learning experiences is through Intelligent Tutoring Systems (ITS). These systems simulate the role of a human tutor by providing tailored feedback and guidance and have been successfully applied in computer-based language learning (e.,g. Slavuj et al., 2015; Heilman and Eskenazi, 2006). ITS leverage a well-established framework that includes a student model, a domain model, and a pedagogical model (Alhabbash et al., 2016). These frameworks can also be adapted to broader adaptive systems, incorporating instructional models and adaptive engines to enhance personalization (Martin et al., 2020).

Building on the principles of ITS, adaptive learning systems have shown particular promise for heterogeneous learner groups (Taylor et al., 2021), such as Kanji learners. These systems adjust parameters based on individual needs, personalizing the learning process through algorithmic changes and fostering self-regulated learning (SRL) without requiring manual customization. Given the complexities of Kanji networks, such personalization could significantly enhance learner engagement and accommodate diverse needs. In language learning, adaptive systems have been studied less frequently than in STEM education. Kaur et al. (2023) reviewed 1,342 articles, finding that most focus on English grammar and vocabulary, with adaptations targeting problem-solving assistance, learning path recommendations, content adjustments, and domain model modifications. Adaptations can also be based on factors such as knowledge level, performance, cognitive abilities, learning style, and behavior. There are specific adaptive learning tools for Japanese, such as an app developed by Ng et al. (2015), designed for Chinese native speakers to learn Japanese vocabulary and writing. The app offered tasks like selecting correct pronunciations from multiple-choice options or practicing writing and adjusted learning paths based on user performance. While a small study reported positive outcomes, the criteria for determining when to adjust the learning path were not clearly defined. Another example is an adaptive email-based program by Li et al. (2009), which customized Kanji quizzes based on user interests, performance, and preferred times. The system dynamically adjusted test difficulty and sent quizzes at individual times to encourage regular learning. Although the adaptive group in a small study performed better, limited user engagement reduced the system’s overall effectiveness. These studies demonstrate both the potential and the challenges of applying adaptive methods in Japanese learning.

In summary, while existing methods address some needs of Kanji learners, there remains a need for a tool that effectively bridges the gap between network visualization, adaptive learning, and practical memorization techniques. Building on the insight of Zimmerman and McMeekin (2020), it is more practical to integrate such a well-established tool like Anki into a unified approach, combining its strengths with adaptivity and etymological insights to provide comprehensive support for Kanji learners.

4 AN ETYMOLOGY-DRIVEN ADAPTIVE KANJI LEARNING TOOL

This section presents the design decisions for a prototype that addresses the research gap identified in the previous section. The goal is to develop a tool for Jōyō Kanji—characters used in everyday Japanese—that supports mnemonic creation through etymological insights and facilitate self-regulated learning (cf. Zimmerman, 2000). Hence, the tool should be tailored to accommodate diverse learning strategies and adapt to users' varying knowledge levels. It should encourage exploratory learning by allowing users to search for Kanji, explore related characters, and save them for future review. An important consideration is how to effectively present the interconnectedness of Kanji, ensuring that learners can easily navigate relationships between characters. Learners should be able to track their progress by viewing previously studied Kanji and receive personalized learning recommendations based on their known characters. Additionally, the tool should offer simple example sentences and highlights Kanji that have not yet been mastered. It is not the goal, however, to build yet another rote memorization or flashcard learning tool (cf. Zimmerman and McMeekin, 2020) but to allow explorative Kanji learning and to provide an interface to the established and widely used flashcard tool Anki.

The prototype offers three main functions: Looking up and displaying *all* Kanji and their relationships in a visual format, getting detailed information relevant for learning on selected Kanji, exporting these information to the flashcard software Anki, and importing previously learned Kanji via Anki. In addition, the application reacts adaptively to the user's level of knowledge by making learning path recommendations and providing simple example sentences generated with GPT-4 Turbo. The prototype was developed as a platform-independent single-page web application and uses the MEN stack (MongoDB, Express.js, Node.js). It offers various ways to look up Kanji. Users can enter a Kanji using a digital Japanese keyboard or draw a Kanji using their mouse/finger. This function uses Kanjicanvas.js (Klein, 2021), which works even with imprecise strokes and displays a selection of recognized Kanji from which the desired one can be selected.

A central component and the entry point of the tool is an interactive Kanji graph. This graph includes the Jōyō Kanji, their graphemes, and visualizes the construction of the Kanji as a network by using its graphemes as nodes and its etymological composition as edges. The graph is fully interactive, supporting click, drag, and zoom functionalities. The Kanji network can be viewed as a complete graph with a concentric layout, placing Kanji most often used (i. e., the graphemes) in the center. This gives learners the opportunity to understand that a few graphemes build most of the Kanji and are therefore useful to learn to gain a better understanding of many Kanji. This graph addresses the gap that existing tools often fail to visualize the etymological relationships and structural composition of Kanji in an accessible and interactive manner, leaving learners without a clear understanding of how graphemes and other Kanji connect to the Kanji they form. The Kanji network is based on a custom database developed specifically for the prototype, compiled from openly available resources, including (Tamaoka et al., 2017) and CHISE.org. CHISE.org provides a grapheme-based decomposition of Kanji, identifying meaningful structural units rather than merely the smallest components. However, it serves primarily as a character processing database and does not structure this information for learning purposes. (Tamaoka et al., 2017), in contrast, offers an etymological categorization of Kanji but does not include any decomposition into components. Since neither source alone fully supports a pedagogically guided decomposition, our approach refines and integrates these resources to ensure that Kanji are broken down into meaningful and instructive graphemic units, preserving relevance for learners.

The overall design of the tool is illustrated in Figure 3. At the top center (1), users can search for Kanji using the search bar. The central area features the Kanji graph (2), which displays either the entire set of Kanji or a filtered subset along with their relationships. As detailed in Section 4.1, the graph adaptively highlights nodes corresponding to Kanji already known to the user. On the lower left (3), a legend provides a clear overview of the various node and edge types based on the etymological categories. The legend is visible by default but can be toggled as needed. On the right (4), an information box, referred to as the "Kanji card", is displayed whenever a Kanji is selected, offering additional details about the character.

Upon searching for a specific Kanji (see Figure 3), the tool presents a tailored, detailed hierarchical graph

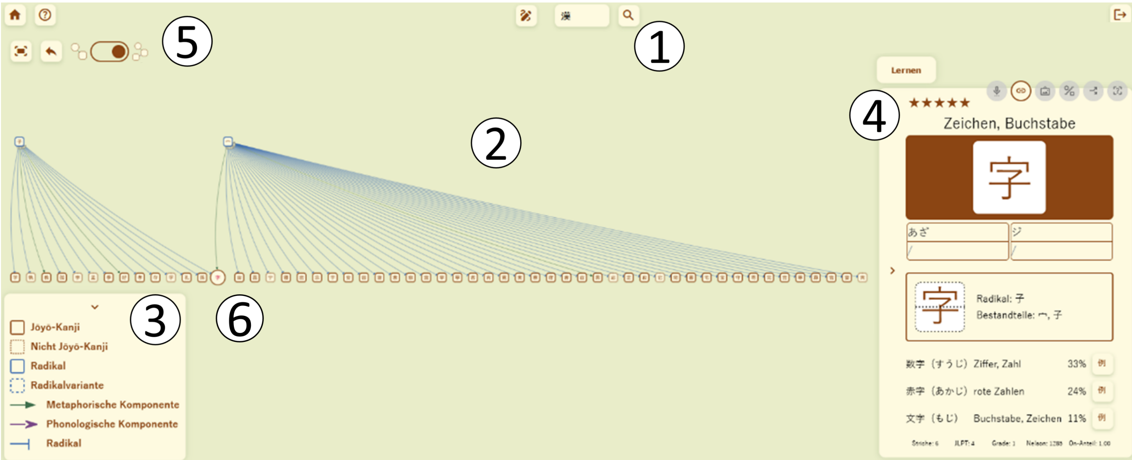

Figure 3: The prototype's user interface after searching for the Kanji 字; details are in the text.

of its parent and child characters. It shows the Kanji together with its direct and optionally also its indirect neighbors. The user can switch between these views to either reduce the number of shown Kanji or get a better overview using the toggle button (5). The layout is based on the Klay Layered algorithm (Schulze et al., 2014), which arranges the nodes according to the hierarchy of the Kanji's graphemes. The selected node is highlighted (6). A special feature of the detail graph are the compound nodes, which represent the relationship between radicals and their variants by grouping radical variants under a common parent node. This structure was chosen to make it easier to understand the composition of Kanji and their variants. The graphs are created using the Cytoscape.js library (Franz et al., 2015), with various optimizations applied to improve performance and reduce complexity.

The Kanji card (see Figure 4) presents detailed Kanji information in an interactive, flashcard-like format. Additional details are revealed through hover interactions. At the top left (1), the usage frequency of the Kanji is displayed (based on Tamaoka et al. (2017)). To make the distribution of the Kanji easier to understand, it is divided into quintiles and visualized using one to five stars. Other important information on the card includes the etymology category of the Kanji, which is represented by icons (2). Two meanings (3) of the Kanji in German are generated using GPT-4, queried with a compact prompt and stored in the database. GPT-4 was chosen over standard dictionary resources due to its flexibility in generating concise, user-friendly meanings by providing only the most important meaning of a Kanji rather than listing all possible interpretations. This approach ensures a consistent format and seamless integration into the application while minimizing the manual effort required for curating definitions. Below this, the Kanji is displayed in large letters (4), followed by a tabular overview of the Japanese readings (5). In the center (6) of the card is the structure display, which shows either the ancient writing of the Kanji for clearer pictorial representation or its composition of graphemes. Further down, the card contains example words and sentences (7), also generated by GPT-4. These example sentences are generated to provide clear context for how the Kanji is used in practice, utilizing grammatically correct and vocabulary-appropriate sentences at a simple language level. The example sentences were generated, when the user clicked on the corresponding button. At the bottom is general data (8) such as stroke count and the level on the Japanese Language Proficiency Test (JLPT).

While the Kanji card allows learners to delve deeply into the details of individual characters, this represents only one facet of the prototype's capabilities. The tool is also designed to encourage broader exploratory learning by adapting dynamically to the user's knowledge level and offering personalized recommendations. These adaptive functionalities, alongside the seamless integration with external tools like Anki, will be elaborated upon in the following section.

4.1 Adaptivity

The prototype adapts learning content to each user's knowledge level, following the adaptive learning framework by Martin et al. (2020). In the prototype, the *Domain Model* is represented by the database containing the Jōyō Kanji, their meanings, and etymologies. The *Student Model* contains the set of Kanji that the learner "knows", and the *Instructional Model* consists of algorithms for learning recommendations.

A central component of the prototype is its integration with the flashcard software Anki, enabling the

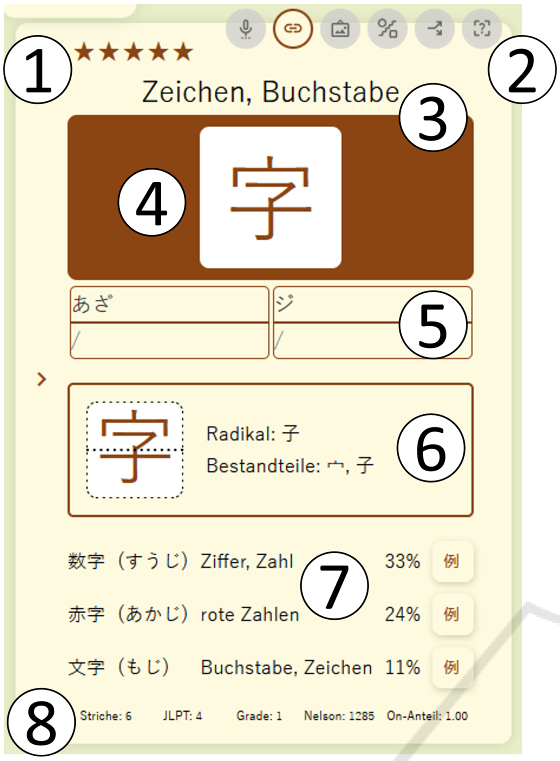

Figure 4: The information card for the Kanji 字.

prototype to rely on already existing data about the user's learning progress. By utilizing Anki's *Ease Factor*, a measure indicating how well a Kanji has been learned, the prototype avoids the need to collect data from scratch, thereby mitigating a potential cold start problem. This integration allows the prototype to dynamically adjust its learning recommendations and visualizations based on the user's current knowledge.

The adaptive nature of the prototype stems from its ability to import knowledge levels of Kanji the user has already learned, as well as to export Kanji for further practice in Anki. This dual functionality ensures a seamless connection between the learning tool and the user's established study habits, enhancing both flexibility and personalization. Details regarding the technical implementation of this integration are provided in the following.

To assess the user's knowledge level, the AnkiConnect[9] plugin is used, which provides a HTTP API allowing external applications on the same computer such as webbrowsers to exchange data with Anki. The prototype makes use of this API to allow users to export Kanji they want to learn as flashcards to Anki or to import the state of knowledge of Kanji they have already learned. However, integrating this functionality into a website requires the domain to be manually added to an allow list in the plugin settings to configure cross-origin resource sharing (CORS) headers

When exporting Kanji to Anki, a pop-up dialog allows the user to select which information (e. g., meaning, etymological details) should appear on the Anki flashcard. The flashcards are automatically generated using a predefined HTML and CSS layout. For importing Kanji—used to update the *Student Model*—the prototype displays a list of Anki decks (folders) from which the user can select. Kanji from these decks, along with their corresponding knowledge levels, are then imported into the prototype.

The prototype visualizes the user's knowledge in the interactive graph by highlighting learned Kanji. Additionally, a special view displays all learned Kanji in a grid layout, sorted by Anki's *Ease Factor*. Kanji considered to be poorly learned are marked in red. The prototype suggests which of these Kanji should be reviewed and encourages users to create personalized mnemonics and adjust the flashcards in Anki.

4.1.1 Learning Recommendations

The prototype generates learning recommendations based on the user's existing Kanji knowledge, promoting an exploratory approach to learning. Once a Kanji is exported to Anki, the application can suggest another Kanji that aligns with the user's learning progress and encourages further discovery. These recommendations leverage the structural relationships between Kanji, prioritizing content that is both relevant and conducive to deepening the learner's understanding through exploration.

There are four main categories of learning recommendations in descending importance:

Recommendations based on radicals: Radicals and their common variations are prioritized, as they appear in many Kanji. The prototype's database includes 208 radicals and 36 variations, with the 25 most frequent radicals appearing in about 61 % of all Kanji. Hence, mastery of these radicals is considered a foundation for understanding a majority of Kanji.

Phonological graphemes: Kanji containing phonological graphemes that influence the pronunciation of other Kanji are recommended if they are deemed relevant for daily use. Relevance is determined using factors identified by Toyoda et al. (2013), who assign scores to Kanji based on these criteria. A grapheme is considered useful by the prototype if its combined score exceeds the average across all factors.

Kanji in example words presented on the Kanji card: If a user selects a sample word containing an unknown Kanji during export, that Kanji is recom-

[9] https://foosoft.net/projects/anki-connect/, last accessed: 2025-02-11

mended for further study, reinforcing learning through contextual connections.

Phonological usage: If the exported Kanji is used as a phonological grapheme in other Kanji, one of these Kanji is recommended for further learning to strengthen pronunciation connections. Learning another Kanji with the same grapheme and pronunciation can enhance recall.

Through these adaptive learning recommendations and the visualization of knowledge, the prototype provides a personalized learning experience tailored to the user's individual progress.

The prototype addresses key use cases, such as repeatedly searching for Kanji, exploring the component-based structure, contextualizing known Kanji with new characters, and using the provided resources for both educational and personal interest.

5 EVALUATION

The goal of the evaluation was to understand how learners perceive and use the tool's visualization and recommendation features in a realistic setting, where they use the tool to study Kanji, and to determine to what extent these features motivate and engage them in exploring and deeply learning Kanji.

To ensure that the tool incorporated both pedagogical insights and features that support self-regulated learning, feedback was sought from a Kanji learning expert for non-native speakers before the evaluation took place. During an hour-long session, the expert reviewed the prototype and suggested improvements, including stroke order animations and a handwritten-style font. She praised the integration of GPT, the graphical representation of Kanji, and the Anki connection as valuable features that promote a comprehensive and flexible learning experience. The expert emphasized the importance of self-regulated learning, particularly for advanced learners, highlighting that the tool effectively encourages users to actively think and make decisions, which she deemed essential for meaningful learning.

The developed prototype was investigated using a user study in May 2024 with 19 participants. The evaluation methodology involved a field study with two versions of the Kanji learning tool: one with adaptive features (the full version as described in Section 4) and one without. The non-adaptive version was visually indistinguishable, but does not provide the buttons to export and import Kanji from and to Anki.

To minimize confounding variables like prior Kanji knowledge or learning strategies, a within-subject design was used: All participants tested both versions in randomized order over the course of two weeks with one week for each version. At the end of each week the participants filled out a questionnaire. The usage of the prototype and which Kanji to look up were not pre-specified, but instead the participants were asked to integrate the tool into their already existing learning routine. Usability and user experience were evaluated using the System Usability Scale (SUS; Brooke, 1996) and Likert scales, which measured satisfaction with features such as the graph, sentence examples, and open-ended questions. Participants were also asked to state their learning goal for the week and rate their motivation to achieve it on a Likert scale. Participants' responses were standardized to a scale from 0 to 4, where 0 represent the most negative rating and 4 the most positive. The activity of the participants in the prototype was logged by registering their clicks.

All 19 participants completed both week's surveys. Thirteen had formal Japanese education, while six were self-taught. Recruitment was through personal connections and snowball sampling, with all participants being German-speaking young adults. The median Kanji knowledge was 350, with five knowing over 1,000, covering a diverse learner range.

Non-parametric statistical tests, particularly the Wilcoxon-Mann-Whitney and Wilcoxon Signed-Rank tests, were applied to compare groups and conditions, and Spearman's rank correlation (ρ) was used to examine relationships. Typically, these tests are stricter than their parametric counterparts Kaur and Kumar (2015), but do not require e. g. normal distributed data. The level of significance is $\alpha = .05$. $\tilde{x}$ denotes the median.

5.1 Results

This section presents the results of the study, beginning with overall findings on usability and user experience, followed by an analysis of participant groups based on their use of the tool's adaptive features.

The participants used the prototype for a median of 57 minutes and rated it favorably, with a usability score (SUS) of 72.1, indicating a good user experience (Bangor et al., 2008). Ten participants rated the usability as good, and four as excellent. SUS scores improved over time, rising from 70.1 in the first week to 74.1 in the second week. Thirteen participants indicated they would use the tool regularly ($\tilde{x} = 3.0$), finding it easy to use and its functions well integrated ($\tilde{x} = 3.0$).

Participants rated the presentation, loading time, and interactivity of the overall graph positively ($\tilde{x} = 3.0$). However, the majority found the graph overwhelming, especially when they did not have a specific learning objective ($\tilde{x} = 1.0$). In contrast, the detailed graph received significantly better ratings for presentation, loading time, and interactivity (all $\tilde{x} \geq 3$), and was perceived as less overwhelming ($p < .001$; $r = .05$) (see Figure 6). It was also seen as more motivating and supportive for learning, particularly for participants with clear learning goals and extended usage. The detailed graph was praised for helping users understand Kanji connections ($\tilde{x} = 4.0$) and the linkage of radicals with their variants ($\tilde{x} = 3.5$).

Seven participants used the Anki export function. They rated it as very helpful ($\tilde{x} = 4.0$) and fast ($\tilde{x} = 4.0$). The participants appreciated the ability to create cards with customizable information ($\tilde{x} = 4.0$). Two participants were motivated by the tool to install Anki for the first time and began using it regularly. The other participants did not use the feature, citing existing schemes for creating Anki cards (33.3 %) or a lack of active engagement with Anki (33.3 %). Participants who used the export function received an average of 35.2 learning recommendations, of which they followed 38.9 %. Moreover, participants rated these recommendations as both highly helpful ($\tilde{x} = 4.0$) and motivating ($\tilde{x} = 4.0$). They perceived the recommendations as non-disruptive ($\tilde{x} = 3.5$) and easily comprehensible ($\tilde{x} = 3.5$). The participants rated the import function as motivating ($\tilde{x} = 4.0$), as they can see their progress, and helpful for understanding the relationships between known and "new" Kanji. The highlighting of poorly learned Kanji was perceived as motivating ($\tilde{x} = 3.0$); however, only two participants actively clicked on and revised poorly learned Kanji, as reflected by a low median agreement score of $\tilde{x} = 2.0$ for the statement: "I followed the request to revise poorly learned Kanji."

In the following, an in-depth analysis is conducted and the participants are grouped based on their actual use of the tool's adaptive features: nine participants used the adaptive functions, while ten did not (despite being able to do so). These participants primarily used the tool to look up Kanji as they were free to decide how to integrate the prototype into their learning process. For instance, 12 participants did not export Kanji to Anki or 15 participants did not explore example more than two sentences. The dimensions analyzed included prior Kanji knowledge, learning goal motivation (whether a learning goal was set and how high the motivation was to achieve it; rated 0–10), usage time, SUS scores, and adaptivity usage.

The analysis revealed a strong positive correlation between learning goal motivation and both the time spent using the tool ($\rho = .66$, $p = .002$) and SUS scores ($\rho = .59$, $p = .008$), as well as the adaptivity usage ($\rho = .59$, $p = .007$) (see Figure 5). Additionally, SUS scores were moderately correlated with usage time ($\rho = .49$, $p = .034$) and adaptivity usage ($\rho = .44$, $p = .057$). Notably, participants who set a learning goal rated the tool significantly higher (SUS = 76.9) than those without a goal (SUS = 62.5; $p = .025$). They also rated the detailed graph lower across several factors, as shown in Figure 6.

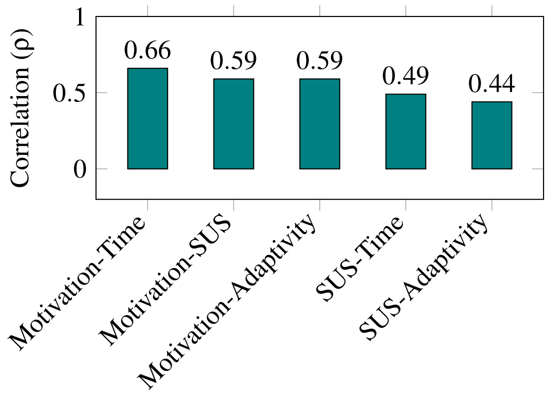

Figure 5: Correlation coefficients (ρ) between key dimensions: learning goal motivation (Motivation), system usability scores (SUS), usage time (Time), and adaptivity usage (Adaptivity).

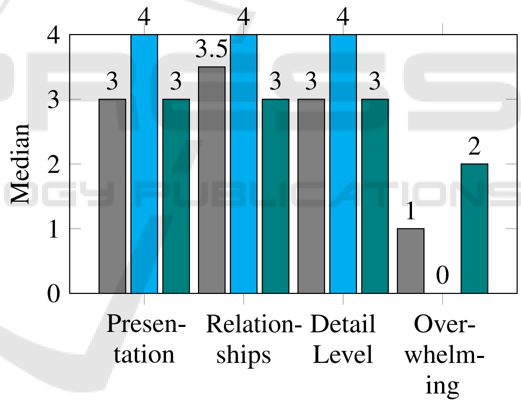

Figure 6: Median ratings for the detailed graph by overall responses (Gray), with learning goal motivation (Cyan), and without learning goal motivation (Teal). Categories: *Presentation* reflects overall satisfaction with the graph's presentation; *Relationships* assesses whether the graph helped identify connections between Kanji; *Detail Level* refers to the appropriateness of the graph's level of detail; *Overhelming* measures whether the graph was overwhelming.

Finally, participants suggested improvements in open-ended comments such as integrating stroke order animations, using a font that looks like handwriting, or providing an introductory video. Many appreciated the tool for combining multiple resources. One participant likened the application to Google Maps, describing how it allowed her to recognize Kanji connections and learn radicals from a "bird's-eye view".

6 DISCUSSION

The study highlights the critical role of learning goal motivation in driving engagement with the prototype. Participants with a clear learning goal spent more time exploring the features and using more adaptive functions. This suggests that internal motivation, when coupled with clear objectives, can amplify the effectiveness of learning tools. Conversely, participants without explicit goals demonstrated less commitment, indicating the need for external incentives to sustain their engagement. The goal of the study was to have an authentic usage by the learners, hence, no clear objectives for tool usage were provided. Having such objectives may have increased tool usage—still, the study provided interesting insights into the design of such a learning tool. To support learners, the prototype may be extended to provide help setting learning objectives (cf. foresight phase in SRL; Zimmerman, 2000) or to include gamification elements such as progress badges, streak rewards, or social competition, which could provide motivate users to engage more actively.

While the SUS score was generally favorable, certain usability challenges emerged. The overall positive rating reflects the intuitive interface and ease of use, particularly for exploration and Anki integration. However, the adaptive features—designed to be a central aspect of the tool—received less emphasis. This suggests that, while users valued adaptability, the core usability features were more immediately impactful. Future research should introduce the adaptive features through onboarding tutorials or specific prompts.

The overall Kanji network graph received mixed feedback. While the majority of users appreciated its interactivity, the sheer volume of information presented was overwhelming for some. Despite this, the inclusion of the overall graph serves an important purpose: it provides learners with a macro-level perspective of Kanji interconnections, enabling them to see how a small subset of foundational Kanji forms the basis for many others. This “big picture” view, though initially daunting, can inspire curiosity and highlight the systematic nature of Kanji learning. Hence, it was initially included. Future refinements could include a personalized knowledge graph to reduce overwhelm and individualize the user experience. This personalized graph could display only the Kanji known to the user together with their parents and children. In contrast, the detailed graph was rated significantly higher for its clarity and its more user-friendly, less overwhelming design. The smaller number of nodes and the use of highlighting for learned Kanji created a more personalized and engaging experience. Participants valued the motivational aspect of seeing their progress through color-coded nodes and exploring etymological relationships in smaller, manageable graphs.

The Anki export function was appreciated by those who used it, as it simplified the process of creating custom flashcards. However, the relatively low adoption rate suggests that the setup effort, particularly in configuring the CORS headers in the Anki HTTP server, posed a barrier for some participants, despite the availability of a tutorial. Simplifying the connection process could address these concerns. Requests for improvement concerned technical aspects such as destination selection and batch export. Although learning recommendations were not frequently used, their quality was rated positively. While only 38.9 % of the recommendations were followed, which may seem low at first glance, it is worth noting that learners typically do not implement every suggestion in tools designed for SRL, especially when presented with a high volume of recommendations. This implementation rate demonstrates a meaningful level of engagement with the tool and reflects the perceived relevance and feasibility of the recommendations.

The import function was highly appreciated, as the users found the display of previously learned Kanji helpful. This feature helped learners recognize their progress and adjust their priorities — both are important aspects in SRL (cf. Zimmerman, 2000). However, some participants suggested extending the functionality to provide methodical learning suggestions for less experienced users. The knowledge level check motivated some participants to adjust their learning priorities, but was not consistently used. Adding beginner-friendly learning strategies could support users in building effective study habits.

7 THREATS TO VALIDITY

While the study provides valuable insights into the usability and impact of the Kanji learning tool, several threats to validity need to be acknowledged.

The sample consisted of only 19 participants. Although efforts were made to include individuals with varying levels of prior Kanji knowledge, all participants shared a similar cultural and linguistic background as German-speaking learners of Japanese. This homogeneity could influence the results, as the tool may perform differently with users from other linguistic or cultural contexts. Future studies should aim to include a larger and more diverse sample to increase the robustness and generalizability of the findings.

Participants tested the tool in their own homes rather than in a controlled environment, allowing for authentic insights into its usability in a real-world context. However, participants were free to use the tool as they wished, without necessarily integrating it into their study routines as suggested. This lack of oversight may have introduced inconsistencies that could affect the validity of the findings.

To minimize bias, the order in which participants used the adaptive and non-adaptive versions of the tool was randomized. However, it is still possible that the novelty or frustration experienced in the first week could have subtly influenced user behavior, perceptions, or exploration in the second week.

Although the study intentionally included participants with varying levels of Kanji knowledge, their existing familiarity with Kanji and learning tools likely influenced their interaction with the prototype. For instance, more experienced learners may have found certain features redundant, while beginners may have struggled to fully utilize advanced functionalities. For example, a button was labeled with 例 ("example") as done in many books to indicate that clicking it would provide an example. However, this label was not easily understood by all beginners. These differences could have affected their usability ratings and engagement levels, highlighting the need for segmentation of user feedback based on experience levels in future studies.

The study was conducted over a relatively short period of two weeks. While this time frame allowed for initial insights into user engagement and satisfaction, it was not sufficient to evaluate the tool's impact on sustained learning outcomes. A longer evaluation period would be necessary to get a more comprehensive understanding of the tool's effectiveness and usability.

8 CONCLUSIONS AND OUTLOOK

In this paper, an interactive learning tool for Kanji was proposed, combining graph-based visualizations and adaptive features to support self-regulated and exploratory learning. The tool enabled learners to explore relationships between Kanji, track progress, and connect new concepts to prior knowledge, while adaptive recommendations aligned with their knowledge levels. The findings underscore the importance of balancing comprehensive features with a user-centric design. Features like the overall graph, while ambitious in scope, need to be complemented by mechanisms that simplify and personalize the user experience. Adaptive visualizations, streamlined integrations, and guided tutorials are critical for accommodating diverse user needs and knowledge levels.

In terms of adaptivity, the tool highlighted the value of context-sensitive recommendations of Kanji aligned with learners' knowledge levels, while preserving the autonomy to engage with or bypass these suggestions. Future developments could benefit from implementing adaptive features in a subtle, non-intrusive manner, catering to learners who prefer flexibility and minimizing disruptions to established learning strategies. The visualization approach and adaptive principles are transferable to other domains, such as exploring German word formations with prefixes and suffixes, visualizing chemical compounds and their functional groups, or illustrating relationships in mathematical formulas and transformations.

The prototype also supports encoding and retrieval; this includes integration with an established tool focused on retrieval, demonstrating that combining specialized tools can enhance the overall learning process. Future tools for supporting memorization may benefit from similar partnerships rather than attempting to create all-encompassing systems. Open Science plays a crucial role in enabling such applications, preventing redundancy and fostering collaboration by allowing researchers to build on existing frameworks rather than reinventing tools and methodologies. The project website of KanjiCompass is https://www.tel.ifi.lmu.de/software/kanjicompass/, where also the source code is available as open source.

Finally, motivation played a key role in tool usage. Learners with clear goals found the tool beneficial for deepening their knowledge, while those lacking intrinsic motivation were less likely to engage fully. Future studies could explore how extrinsic motivation techniques—such as gamification or structured integration into formal learning settings—might increase engagement, especially among learners without predefined goals. For these learners, supporting SRL, particularly in the goal-setting and planning phases, is essential. Future developments should focus on further promoting learner motivation and strengthening SRL.

Finally, the study highlighted the potential of large language models (LLMs) in generating example sentences and translations. Future research should investigate their capabilities and quality further, exploring how they can enhance learning tools across domains.

REFERENCES

Alhabbash, M. I., Mahdi, A. O., and Naser, S. S. A. (2016). An intelligent tutoring system for teaching grammar english tenses. *European Academic Research*, 4(9):1–15.

Bangor, A., Kortum, P. T., and Miller, J. T. (2008). An empirical evaluation of the system usability scale. *Int. J. Hum.-Comput. Interact.*, 24(6):574–594.

Banno, E., Ikeda, Y., Shinagawa, C., Tajima, K., and Takashiki, K. (2009). *Kanji Look and Learn: 512 Kanji* *with Illustrations and Mnemonic Hints - イメージで覚える［げんき］な漢字512*. The Japan Times, 5th edition.

Brooke, J. (1996). Sus - a quick and dirty usability scale. In Thomas, P. W. J. B., Weerdmeester, B., and McClelland, I. L., editors, *Usability Evaluation In Industry*, pages 189–194. Taylor & Francis Ltd.

Cook, V. (2016). *Second Language Learning and Language Teaching: Fifth Edition*. Routledge, 5 edition.

Crowley, D. P. (1968). The occurrence and function of chinese characters in modern japanese orthography. *The Journal-Newsletter of the Association of Teachers of Japanese*, 5(3):1–9.

Franz, M., Lopes, C. T., Huck, G., Dong, Y., Sumer, O., and Bader, G. D. (2015). Cytoscape.js: a graph theory library for visualisation and analysis. *Bioinformatics*, 32(2):309–311.

Hanson, A. E. S. and Brown, C. M. (2020). Enhancing l2 learning through a mobile assisted spaced-repetition tool: an effective but bitter pill? *Computer Assisted Language Learning*, 33(1-2):133–155.

Heilman, M. and Eskenazi, M. (2006). Language learning: Challenges for intelligent tutoring systems. In *Proc. Workshop of Intelligent Tutoring Systems for Ill-Defined Tutoring Systems.*, pages 20–28.

Heisig, J. W. (1994). *Remembering the Kanji I: A complete course on how not to forget the Meaning and Writing of Japanese characters*. Japan Publications Trading Co., Ltd., 3rd edition. 12th printing.

Jeronimus, M., Westerveld, S., van Leeuwen, C., Bhulai, S., and van den Berg, D. (2017). Japanese kanji characters are small-world connected through shared components. In Bhulai, S. and Kardaras, D., editors, *DATA ANALYTICS 2017 : The Sixth International Conference on Data Analytics*, pages 53–58. IARIA.

Kaur, A. and Kumar, R. (2015). Comparative analysis of parametric and non-parametric tests. *Journal of computer and mathematical sciences*, 6(6):336–342.

Kaur, P., Kumar, H., and Kaushal, S. (2023). Technology-assisted language learning adaptive systems: A comprehensive review. *International Journal of Cognitive Computing in Engineering*, 4:301–313.

Klein, D. (2021). The stroke correspondence problem, revisited. *arXiv preprint arXiv:1909.11995v3*.

Komarek, A., Pavlik, J., and Sobeslav, V. (2015). Network visualization survey. In Núñez, M., Nguyen, N. T., Camacho, D., and Trawiński, B., editors, *Computational Collective Intelligence*, pages 275–284, Cham. Springer International Publishing.

Li, J. and Zhou, J. (2007). Chinese character structure analysis based on complex networks. *Physica A: Statistical Mechanics and its Applications*, 380:629–638.

Li, M., Ogata, H., Hashimoto, S., and Yano, Y. (2009). Adaptive kanji learning using mobile-based email. In *Proceedings of the 17th International Conference on Computers in Education, ICCE 2009*, pages 520–526.

Martin, F., Chen, Y., Moore, R. L., and Westine, C. D. (2020). Systematic review of adaptive learning research designs, context, strategies, and technologies from 2009 to 2018. *Education Tech Research Dev*, 68:1903–1929.

Mori, Y. (2014). Review of recent research on kanji processing, learning, and instruction. *Japanese Language and Literature*, 48(2):403–430.

Ng, S. C., Lui, A. K., and Wong, Y. K. (2015). An adaptive mobile learning application for beginners to learn fundamental japanese language. In Li, K. C., Wong, T.-L., Cheung, S. K. S., Lam, J., and Ng, K. K., editors, *Technology in Education. Transforming Educational Practices with Technology*, pages 20–32, Berlin, Heidelberg. Springer Berlin Heidelberg.

Rose, H. (2017). *The Japanese Writing System: Challenges, Strategies and Self-regulation for Learning Kanji*, volume 116. Multilingual Matters.

Schulze, C. D., Spönemann, M., and von Hanxleden, R. (2014). Drawing layered graphs with port constraints. *Journal of Visual Languages & Computing*, 25(2):89–106.

Slavuj, V., Kovačić, B., and Jugo, I. (2015). Intelligent tutoring systems for language learning. In *2015 38th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, pages 814–819.

Tamaoka, K., Makioka, S., Sanders, S., and Verdonschot, R. G. (2017). www.kanjidatabase.com: a new interactive online database for psychological and linguistic research on japanese kanji and their compound words. *Psychological research*, 81:696–708.

Taylor, D. L., Yeung, M., and Bashet, A. (2021). Personalized and adaptive learning. *Innovative learning environments in STEM higher education: Opportunities, Challenges, and Looking Forward*, pages 17–34.

The Japan Foundation (2023). Survey report on japanese-language education abroad 2021. Technical report, The Japan Foundation.

Toyoda, E., Firdaus, A. M., and Kano, C. (2013). Identifying useful phonetic components of kanji for learners of japanese. *Japanese Language and Literature*, pages 235–272.

Wozniak, P. A. (1990). Optimization of learning. Master's thesis, Poznan University of Technology, Poland.

Yan, X., Fan, Y., Di, Z., Havlin, S., and Wu, J. (2013). Efficient learning strategy of chinese characters based on network approach. *PloS one*, 8(8):e69745.

Zimmerman, B. (2000). Attaining self-regulation: A social cognitive perspective. *Self-regulation: Theory, research, and applications/Academic*.

Zimmerman, E. and McMeekin, A. (2020). A review of japanese call for kanji, vocabulary, and reading: Findings, best practices and future directions. *JNCOLCTL*, 29:34–68.