

Visualization of Uncertainty in Lattices to Support Decision-Making

Christopher Collins,¹ Sheelagh Carpendale,² and Gerald Penn¹

¹University of Toronto, ²University of Calgary

Abstract

Lattice graphs are used as underlying data structures in many statistical processing systems, including natural language processing. Lattices compactly represent multiple possible outputs and are usually hidden from users. We present a novel visualization intended to reveal the uncertainty and variability inherent in statistically-derived lattice structures. Applications such as machine translation and automated speech recognition typically present users with a best-guess about the appropriate output, with apparent complete confidence. Through case studies we show how our visualization uses a hybrid layout along with varying transparency, colour, and size to reveal the lattice structure, expose the inherent uncertainty in statistical processing, and help users make better-informed decisions about statistically-derived outputs.

Categories and Subject Descriptors (according to ACM CCS): H.5.2 [Information Interfaces and Presentation]: User Interfaces – Graphical user interfaces (GUI); E.1 [Data Structures]: Graphs and networks; I.2.7 [Natural Language Processing]: Machine translation, Speech recognition and synthesis

Keywords: lattice, uncertainty, information visualization, machine translation, speech recognition

1. Introduction

In this paper we present a generalizable decision support visualization that reveals uncertainty in lattices generated by statistical algorithms. While lattice structures are used as black boxes in many processing systems, we know of no other visualization to support their use in “human-in-the-loop” decision making. Our low-cost visualization of uncertainty uses techniques based on well-known properties of human perception. We present case studies applying our visualization to reveal uncertainty in machine translation and automated speech recognition outputs. By supporting exploration of the alternatives considered by these statistical algorithms, our visualization leads to the discovery of better solutions.

While general graphs and some subsets of graphs such as trees have received considerable visualization attention, other important subsets such as lattices have been largely ignored. Lattice graphs are used as the underlying data structures in many statistical processing systems and serve well in holding the possible ranked alternative solutions. Our lattice graph visualization uses a combination of grid and force-

based techniques to create a layout that focuses on multiple encodings of the uncertainties using position, transparency, and colour (see Figure 1). These statistically-derived lattices are amenable to visualization since the uncertainties are locally constrained. Our visualization is readily extensible in general for creating representations of uncertainty in general lattice graphs.

Many current applications, such as the majority of those in the realm of natural language processing [Koe04, Jel98, e.g.], are statistically based. Their outputs represent a “best guess” by the algorithm, given some training data, parameter settings, and input. These best-guess outputs come from a very large collection of possibilities, each ranked with a score. However, these systems present their result in a black-box fashion, showing only a single response. Since no details about probabilities, uncertainties, or the workings of the algorithms are provided, it is easy to misconstrue the output as having a low uncertainty. This lack of detail deprives us of the context necessary to make well-informed decisions based on the reliability of that output.

Understanding about the human reasoning process in-

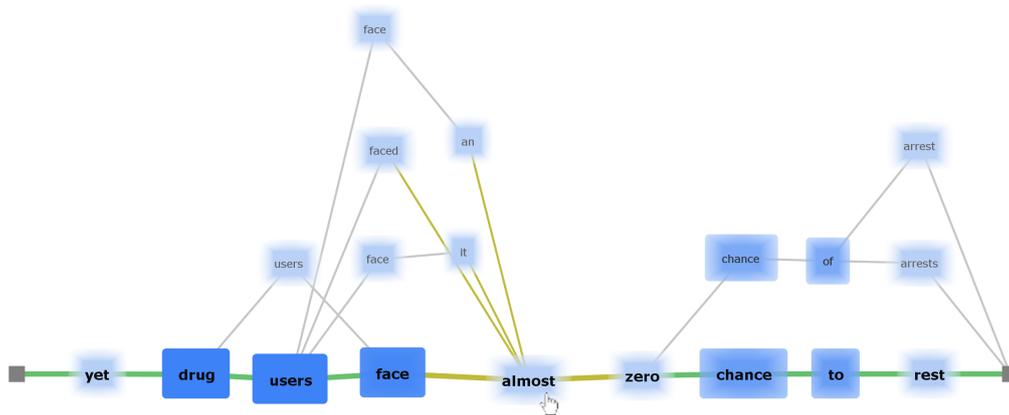


Figure 1: A speech recognition lattice for which the statistically-identified best path is not the best transcription. The best path according to the model is shown in green. Node uncertainty is represented by position, fill hue, and edge transparency gradient. Edges attached to the node under the mouse pointer are coloured gold.

forms us that, while not idealized Bayesian decision-makers, people do make decisions based on their analysis of the objective context of the problem and subjective probabilities informed by their personal body of knowledge [Coh79]. For example, in the context of a natural language system such as machine translation, a person makes a decision about the reasonableness of the output based on their prior knowledge of likely constructs in the language. Based on Cohen’s review of research on reasoning, we work with the assertion that the quality of the decision about whether to accept the algorithm’s best guess can be enhanced by knowing the uncertainty inherent in the solution. That is, providing easy access to the objective context will enable people to make better decisions. Since the effort a person will want to expend in making a decision is proportional to the perceived importance of that decision, we want to keep the algorithm’s best guess obvious while providing visual access to ranked probabilities. For instance, a person may accept a confusing translation in casual conversation in an Internet chat-room, but would reject the same problematic translation in a multi-lingual business negotiation.

There are many aspects of language modelling that statistical processing has yet to master — for instance, an output of speech recognition occurring in the corpus we use is, “the prisoners resisted a rest.” Without our visualization one would not know that “the prisoners resisted arrest” was the second-highest scored hypothesis. While any native speaker would guess the correct reading of the phrase, presenting it visually in parallel with the algorithm’s best guess removes the puzzle aspect for a native speaker but provides a learner with the needed support. By revealing alternative hypotheses considered by the algorithm, and the uncertainties associated with each, our visualization shows promise for facilitating the process of recognizing and correcting of errors.

2. Background

As information visualization as a field has matured, focus on visualizing uncertainty in a dataset in conjunction with the absolute data values has increased [JS03]. Amar and Stasko call for bridging of analytic gaps between data representation and analysis, and one technique they suggest is to expose “black box” processes through visualization of meta-data such as uncertainty [AS04]. Examples from the literature that are relevant to our approach include using line thickness and transparency to represent uncertainty in architectural drawings of ancient structures [SMI99] and using iso-surfaces and motion blur to represent uncertainty in molecular (node-link) diagrams [RJ99]. Zuk and Carpendale [ZC06] present a theoretical analysis of these and other uncertainty visualizations in which they summarize the significant theories of Bertin [Ber83], Tufte [Tuf01], and Ware [War04] and apply them as heuristics for evaluation of visualizations of uncertainty. We draw upon their analysis for design guidance. We will reflect more upon our design choices based on these theories in the following section, after a brief review of lattice graphs and lattice graph visualization.

Formally, a partially ordered set L is a lattice if, for all elements x and y of L , the set $\{x, y\}$ has both a least upper bound in L and a greatest lower bound in L . Lattices are used in formal concept analysis (Galois lattices), and have been previously visualized using simple force-directed layouts [VGRH03]. Lattice drawing has also been of interest to the universal algebra and graph drawing communities, where the focus has been on drawing Hasse diagrams: the edges must be straight lines and the vertical coordinate must agree with the order. Reducing edge crossings has been a primary concern [Fre04]. Our goal differs in that we are not focussed on understanding the particular formal structure of the lattice, but rather using that structure to support understanding of the data and uncertainty represented by it.

The “lattices” in statistical processing do not meet all conditions of this formal definition. Intuitively, we can imagine a lattice in this work as a partial order with a unique beginning and end. Seen as a graph, for every node in a lattice there exists a path from the beginning to the end which passes through that node. To our knowledge, neither lattices for statistical processing nor uncertainty within lattices have been previously visualized.

3. Lattice Uncertainty Visualization

Traditional statistical processing systems use a large corpus of data to quickly produce a single hypothesis, drawing on a computer’s strength in dealing with large amounts of data with the goal of quickly solving a problem. However, if one is presented with the best result of a statistical process, but the quality is so poor it is not useful, then the original goal of providing convenience is not met. Building on the generalization of human-computer optimization by Scott et al. [SLK02], we hypothesize that by including a human “in-the-loop” we can leverage the intelligence of the human and the processing power of the computer to quickly solve the same problems with better solutions.

To meet this goal, we identified several constraints to guide our design process:

- ensure easy readability of paths in the lattice;
- provide an intuitive visual mapping of uncertainty within the lattice which supports the ordering of nodes;
- provide for visual pop-out of nodes of high certainty and nodes in the optimal path identified by the algorithm;
- provide alternative representations of the data to clarify meaning, where possible;
- in most cases, require no interaction;
- where interaction is necessary (providing detail in context and manipulation of best-path tracing), it should be lightweight and easy to learn.

In order to ground our design in an understanding of human perceptual capabilities, we investigated the properties of visual variables [Ber83], leading us to select those that allow for high-speed estimation [HBE96] to convey relevance, and that provide an ordered reading to convey uncertainty. From this, we chose edge sharpness, hue, and position to make nodes of high confidence stand out, and uncertain nodes less visually prevalent. Also, since value, size, position, and transparency are ordered (values can be visually sorted), we used these to encode uncertainty to allow for comparison of the relative scores between nodes.

3.1. Data

The lattices generated by statistical processing are collections of nodes and edges, each associated with a data value (for example, a word) and a score (for example, an uncertainty). A lattice is generated as a representation of the solution space for a given problem; any path from beginning to

end through the lattice represents an hypothesis about the solution. However, the lattice may or may not contain the true solution. For example, a speech recognition lattice contains all the potential transcriptions considered by the algorithm. It may not, in fact, contain the correct transcription. Each lattice has a *best path* through it, based on node scores as well as a *true best path* which, while it may not have the best node scores, best matches the true solution. Our goal is to use visualization to provide an opportunity for people to combine the scores with their world knowledge to discover the *true best path* or to reject the entire lattice.

In a lattice generated by a statistical process there may exist a unique start and end node, representing the beginning and end of all possible solutions. If such endpoint nodes do not exist, we create them, extending edges from all starting and ending nodes to the new endpoints. Unique endpoints provide for an easy to locate visual entry-point to reading the lattice. Our visualization algorithm reads lattices from the statistical processing (source) algorithm using either an interface customized to the application, or the HTK Standard Lattice Format (SLF) (<http://htk.eng.cam.ac.uk>). In our current work, we only use labels and scores on nodes. We convert lattices with edge probabilities to have posterior probabilities on the nodes using the SRILM lattice toolkit [Sto02]. Finally, we retrieve the best path through the lattice, according to node scores, either directly from the source algorithm or using the SRILM lattice toolkit.

When we discuss the uncertainty of lattice nodes, we do not strictly mean uncertainty, as might be quantified by an entropic measure, for example, but rather a more application-specific property that emerges from the lattice which reflects the probability that the node is part of the true best path. In particular, node scores are generally relative confidence scores, not true probabilities, and the presentation of several alternatives at any slice in the lattice is more an indication of the number of plausible solutions, rather than of a small margin of preference among those alternatives. The score of a node, nevertheless, can be interpreted as a measure of certainty that the node is the correct choice for its span and appears in the best path.

3.2. Layout

In the graph drawing community, where the lattices are usually representations of an algebra, the convention is to draw the order vertically, from bottom to top [Fre04]. However, in the languages our visualization is designed to support, text-based reading occurs left-to-right. Additionally, temporal flow (as in the flow of speech signals) is usually thought of as occurring left-to-right. So, to support our design goal of easy readability, we align our visualization horizontally to allow for more natural left-to-right tracing of paths.

Our layout algorithm is a hybrid of grid and force-based layouts. Initially, the lattice graph is laid out on a grid, anchored by endpoints which are positioned according to the

length of the algorithmic best path through the lattice. This avoids large gaps or significant overlaps. Horizontal node positioning is derived from the node order in the lattice from beginning to end. Vertical position is assigned to separate nodes covering the same span, ordered by increasing uncertainty bottom to top. Because the algorithmic best path is of most interest, we place it along the bottom, regardless of the individual node scores. This anchors the visualization in the algorithm's best-guess solution and facilitates easy reading of it (see Figure 2A). Position, the strongest visual variable, according to Bertin [Ber83], ensures that the least important nodes (highest uncertainty) appear furthest from central visual focus along the bottom.

The grid layout can sometimes result in overlaps for nodes with lengthy labels and for larger lattices. We automatically zoom out the display to attempt to fit the entire lattice without overlap, but must limit the scale factor to ensure label legibility. To reduce overlap, we adjust the layout using a force-directed energy minimization. An unconstrained force-directed layout alone would create an unpredictable and unordered layout (see Figure 2B). Thus, nodes are anchored by invisible springs to their grid positions, and to each other by springs represented visually as graph edges. Repellent forces are applied between nodes to reduce overlap and the energy minimization is run for several seconds to stabilize the layout. This hybrid layout allows any overlapping nodes to separate while not moving far from their grid-determined position, balancing the need to keep nodes in the rigid layout for easy left-to-right reading and the demand that nodes do not overlap (see Figure 2C).

3.3. Uncertainty Encoding

Uncertainty in the lattice is foremost visualized through the presence of alternative paths through the lattice: more paths can indicate greater uncertainty, depending on the relative scores for the nodes in each path. Uncertainty scores are used to colour the nodes using a range from saturated blue to desaturated gray. However, continuous colour scales generally should be avoided for numerical data (colour perception varies due to several factors, including the size of items) [War04]. To compensate for this, we redundantly encode the scores in the node border using size and transparency. Hue, border size, and outer edge transparency are all linearly related to the uncertainty score on each node.

We present two alternatives for encoding, each with its own advantages (see Figure 3). In the "bubble border" view, the node border varies from a tight solid blue, indicating high confidence, to a transparent, wide, gray border, indicating uncertainty. Large, semi-transparent borders lead to an intuitive reading of uncertainty. In the "gradient border" view, the node border varies from a crisp edge to a gradient leading to complete blending with the background. The gradient border is achieved through a linear blending of full opacity at the node center to variable transparency at the outer edge. This

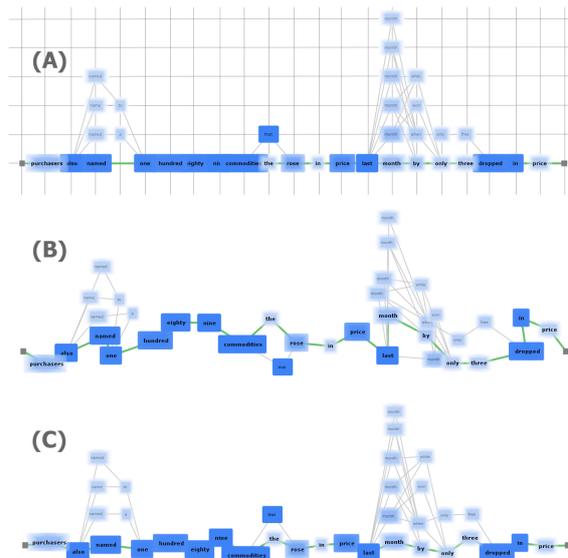


Figure 2: Layout construction: (A) rigid grid-based lattice, (B) force-directed layout, (C) hybrid layout. The hybrid layout provides the regularity benefit of the grid-based layout and the overlap avoidance of the force-directed layout.

effect simulates semantic depth-of-field [KMH01] in which items with crisp focus pop-out. Even though the gradient fill on the nodes in this view does not overlap the text label, in informal testing we found that the blur effect seemed to make the labels more difficult to read. So, while the gradient border may be more intuitive and lead to a more immediate reading, the bubble border may, in the end, be a more usable encoding. In both cases, the use of transparency is supported by visualization theory: transparency blends the visual variables of value and colour in a redundant encoding from which an ordered reading is possible [Ber83]. These techniques satisfy our goal: to coarsely and quickly indicate relative uncertainty without providing specifics on the scores of each node. In fact, the precise numbers are often not very meaningful: they result from the settings of many variable parameters in the model which generated the lattice and are generally only comparative within a particular lattice.

3.4. Interaction

Simple interaction techniques are provided: when hovering over a node, its edges are highlighted in gold to disambiguate edge crossings (see Figure 1). Nodes can also be dragged to further clarify edge attachment, returning to their original location when released. By right-clicking nodes, the user can remove and add nodes to the green-edged best path, thereby using their knowledge of the context (for example, their prior linguistic knowledge) to reassign the best path (see Figure 4). Where others have used iterative feedback to recompute a new best path through a (hidden) lattice based on user feed-

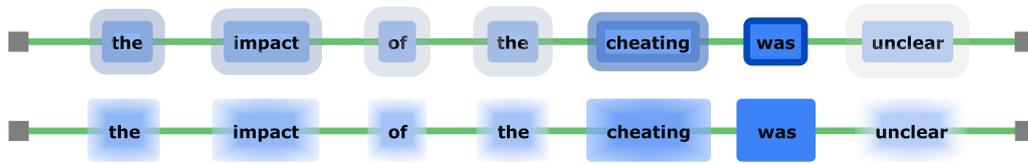


Figure 3: Two alternative encodings of uncertainty. The top, “bubble border”, uses border thickness, hue, and transparency to convey uncertainty: tight, solid blue borders have a higher certainty. The bottom, “gradient border”, uses blurring of the edges through a gradient and variable transparency: more solid borders have higher certainty. Both variations also use hue variation from bright blue (high certainty) to gray-blue (low certainty).

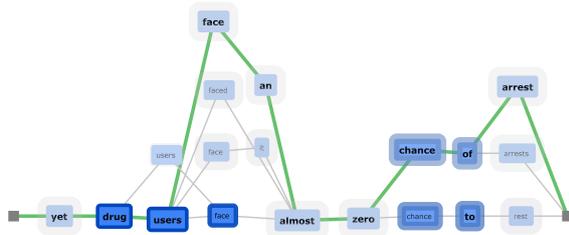


Figure 4: Lattice from Figure 1 with best path corrected.

back [LS06, e.g.], we provide complete control to the human decision maker. For our interface, an iterative process is unnecessary as the entire lattice is visible. Furthermore, iterative interaction would violate our minimal interaction design constraint. In the case studies to follow, we will explore how this functionality can be applied in real implementations.

4. Case Study: Machine Translation

Machine translation offers much promise for improving workplace communication among colleagues situated in offices in different parts of the world. Many corporations use instant messaging chat as a means of facilitating communication, however current translation quality is too low to feasibly use it in a critical setting. In this case study we present a prototype visualization system for instant messaging conversations which uses our lattice uncertainty visualization to reveal the uncertainty in the translation and provide alternative translations when available (see Figure 5).

Despite evidence that social spaces in the Internet are multilingual in nature, these spaces still lack rich cross-linguistic communication [HPRV*07] and little research has been directed toward supporting cross-lingual chat. Recent studies on cross-lingual instant messaging chat in distributed workplaces show that problems with translation quality do negatively effect conversations [YI06]. To our knowledge, only a few commercial cross-lingual chat applications (e.g., <http://www.chattranslator.com>) exist and they only present the best-path solution to the user. The Cairo system [SJ00], a tool for translation researchers, is related to our visualization in that it provides a means for exploring

alternative translations. However, where we focus on providing a visual means to understand translation uncertainty, the Cairo interface is tailored for examining and evaluating specific word correspondences between languages.

4.1. Translation Architecture

We chose to work with instant messages as the data for uncertainty visualization in translation because they offer several advantages for this work. They tend to be short, keeping translation time low and providing an appropriate amount of data for a small-scale visualization. The result should be a manageable number of alternate translations for chat participants to explore. We developed a bidirectional instant messaging client which performs translation on messages it receives using a beam search decoder for statistical phrase-based translation models. We trained the decoder, Phramer [Olt06] (an open-source implementation of [Koe04]), on the the English, Spanish, French, and German portions of the Europarl corpus (approximately 1M sentences in each language) [Koe]. The phrase-based translation is supported by a trigram language model trained on the same data [Sto02]. The translation algorithm evaluates the input data and creates a set of translation hypotheses, assigning confidence scores to each word and phrase based on occurrences in the corpus. The best path through the lattice, according to the scores, is labelled by the translation system. Using this data, we create a compact lattice populated with all alternate translations which have a score within a preset threshold of the best score. This graph, complete with scores for each node, is then used as the lattice for visualization.

4.2. Interface

In following with norms of instant messaging client design, we maintain a chat transcript: the green-edged best path is recorded to the chat history when the next message is received. However, it often occurs that a node along this path has a low confidence score (high uncertainty). The user can explore alternative translations for this span of the sentence, or, if no reasonable alternatives exist, use the chat to request clarification from the author of the original message. When out-of-vocabulary words are encountered, or the translation

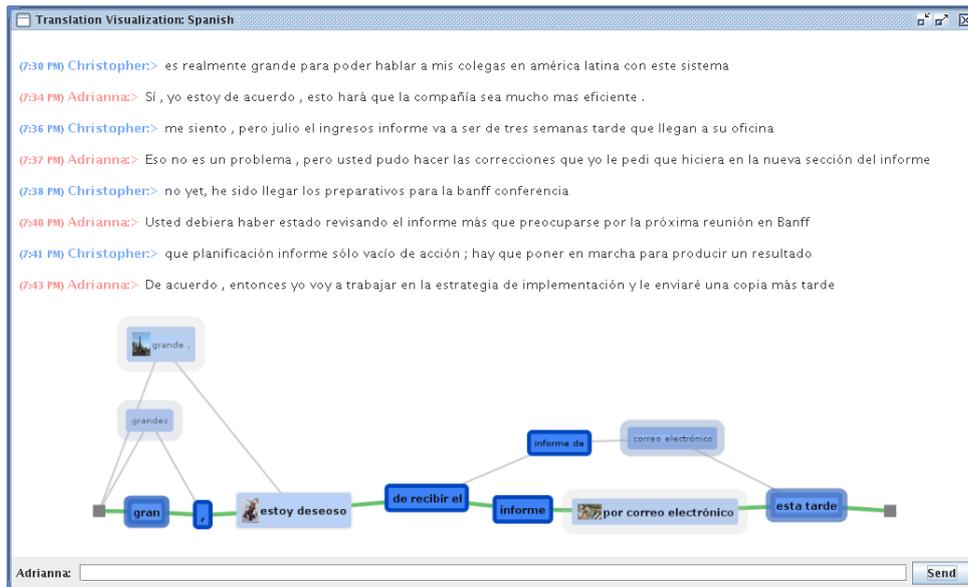


Figure 5: Translation uncertainty visualization in a chat interface: upper panel records chat history in the language of the local conversant, lower panel displays the visualization of the most recently received message. Translations of low confidence are also augmented with representative pictures collected from the Internet.

uncertainty is particularly high, photos are retrieved from Flickr (<http://www.flickr.com>) using the original (untranslated) words as a search query. In some cases, images may easily clarify the intended meaning (see Figure 6).

The main interaction is through the chat message box and reading the data presented in the visualization. To facilitate accurate chat logging, the ability to toggle node inclusion in the green “best path” is provided. In this way, alternate translations can be selected and recorded in the log instead. Selecting a photo node enlarges it, revealing a set of four images about that node.

4.3. Discussion

This chat system was designed for two participants in on-line meeting, each of whom does not speak the other’s language. Through our instant messaging system, they converse, in some cases exploring the lattice uncertainty visualization structure for clarification of a poor translation, and in other cases rejecting the entire translation as too low quality based on the node uncertainties. This visualization and chat system was demonstrated at CSCW 2006 [CP06]. Informal user feedback indicated an interest in multi-lingual chat in general, and in the visualization of uncertainty. Participants indicated they would like to try the system for a longer period of time, in particular they liked the inclusion of photos on untranslatable nodes. From using the visualization, we notice that for English translated to French or Spanish, many of the lattices have ambiguities on single words and short phrases, whereas for English to German there are

longer segments of ambiguity, likely due to the freer word order of German.

5. Case Study: Automated Speech Recognition

Automated speech recognition is another application area where lattices are commonly used in processing but only the best solution is reported. The selection of the best path is dependent on the quality of the speech input signal, the acoustic model, and the language model. With many places to go wrong, speech recognition often produces incorrect results.

There have been investigations into using lattices to suggest alternative translations in drop-down lists and in multi-modal interfaces, including handwriting recognition [SMW01], but generally people remain dissatisfied with these interfaces. Kemp and Schaaf [KS97] report on a text-based tagging system which labels each word of speech recognition output with a normalized measure-of-confidence score. However, in their work, alternative hypotheses are not provided. In all cases, the lattice structure remains hidden from view. Although much attention has been given to supporting correction of transcription errors, we know of none that use the lattice and its scores directly in “human-in-the-loop” interaction.

5.1. Recognition Architecture

Algorithms for automated speech recognition are generally arranged as a pipeline of data transformations. For our purposes, we can think of this pipeline as a three step process:

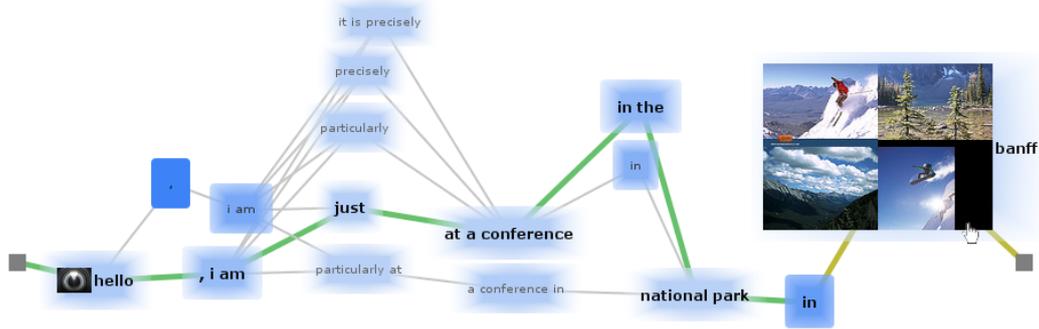


Figure 6: Translation lattice for the German sentence, “Hallo, ich bin gerade auf einer Konferenz im Nationalpark in Banff.” The statistically-identified best path (along the bottom) was incorrect and has been repaired. Photo nodes provide an alternative representation for words not in the translation vocabulary. Mouse over expands the node and reveals four photos, while other nodes move away to avoid occlusion.

(1) an acoustic model takes a digitized speech signal and creates a word lattice with scores, (2) a language model rescores the lattice based on probabilities of words occurring in sequence, (3) the best path through the lattice based on the acoustic and language model scores is output. The NIST '93 HUB-1 collection of word lattices represents data captured from this process after step 2. This collection of 213 lattices comes from high-quality recordings of 123 speakers reading excerpts of the Wall Street Journal. Note that in the HUB-1 collection, some node labels may be repeated, indicating multiple possibilities arising from uncertainty about the start time or length of the word in the speech signal. The lattices include acoustic and language model scores along the edges. We used the SRILM lattice toolkit to calculate scores for the nodes and prune the lattices to contain at most the 50 best unique paths. We also eliminate null nodes (silences) and nodes with scores below 0.01% of the best scoring node. While our visualization is decoupled from the actual speech signal, it could easily be connected to the speech recognition pipeline directly.

5.2. Discussion

Examples of visualization of the HUB-1 lattices appear in Figures 1–4, and there are many examples from this case study for which the best path chosen using the node scores is not the true best path in the lattice. In informal testing, it seemed that in many cases, the correct path was obvious upon reading the optional nodes for a particular span — only one path made sense. Through using the visualization, we discovered that the speech lattices seem to generally have a different structure than translation lattices: where ambiguity in translation often presents an alternative or two for a span of several nodes, speech recognition lattices show highly localized ambiguity (see Figure 7). This stems from the difficulty of acoustic models for speech recognition to recognize short words; a short duration and low signal amplitude lead

to elevated uncertainty. By coupling our visualization of uncertainty with human linguistic knowledge, it is possible to make better informed decisions about the quality of a transcription, and to correct errors by selecting a new path in the lattice. In this way our visualization could support real time editing of speech transcripts on a sentence-by-sentence basis.

6. Directions for Future Work and Conclusion

At the conclusion of our design process, we identified several opportunities for future work on lattice uncertainty visualization. Our visualization relies on embedding of uncertainties on the nodes. Some statistical processing algorithms also provide scores for edges. Extending the visualization to incorporate edge uncertainties is a natural next step.

We would like to conduct user studies to confirm whether our visualization would be preferred over a simple single best solution presented “black box” style, and to what extent it helps people make better decisions about the data. Such a study could be conducted with our the instant messaging client in a multi-lingual distributed workplace. Additionally, multi-modal approaches to correction of speech recognition transcripts have previously been reported, and it would be interesting to discover how interacting with the word lattice directly in our visualization performs in comparison to indirect approaches.

We have presented a generalizable visualization for uncertainty in lattices generated by statistical processing. The techniques for visually encoding uncertainty may be applicable to other node-link structures, such as Hidden Markov Model trellises, probabilistic finite state automata, and general graphs. Following a set of design constraints grounded in a review of relevant literature on human perceptual capabilities and visual variables, we introduce a new hybrid layout which conveys confidence through position while en-

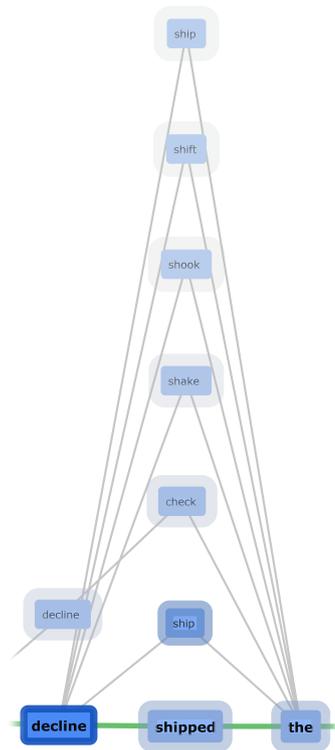


Figure 7: Speech recognition lattices often have localized areas of high uncertainty.

hancing readability of lattices. Our visualization reveals the search space considered by common statistical algorithms in areas such as natural language processing, and could be useful as a teaching tool.

References

[AS04] AMAR R., STASKO J.: A knowledge task-based framework for design and evaluation of information visualizations. In *Proc. of IEEE Symp. on Information Visualization* (Oct. 2004), pp. 143–149.

[Ber83] BERTIN J.: *Semiology of Graphics: Diagrams, Networks, Maps*. University of Wisconsin Press, 1983.

[Coh79] COHEN G.: Subjective probabilities and statistical parameters. In *Uncertain Outcomes*. MTP Press, 1979, pp. 23–48.

[CP06] COLLINS C., PENN G.: Leveraging uncertainty visualization in multilingual chatrooms. In *Proc. of Computer Supported Cooperative Work, Interactive Demonstration Session* (Nov. 2006), ACM Press.

[Fre04] FREESE R.: *Concept Lattices*, vol. 2961 of *Lecture Notes in Computer Science*. Springer-Verlag, 2004, ch. Automated Lattice Drawing, pp. 112–127.

[HBE96] HEALEY C. G., BOOTH K. S., ENNS J. T.: High-speed visual estimation using preattentive processing. *ACM Transactions on Computer-Human Interaction* 3, 2 (1996), 107–135.

[HPRV*07] HERRING S. C., PAOLILLO J. C., RAMOS-VIELBA I., KOUPER I., WRITE E., STOERGER S., SCHEIDT L. A.,

CLARK B.: Language networks on LiveJournal. In *Proc. of the 40th Hawai'i International Conference on System Sciences* (2007), IEEE Press.

[Jel98] JELINEK F.: *Statistical Methods for Speech Recognition*. MIT Press, 1998.

[JS03] JOHNSON C. R., SANDERSON A. R.: A next step: Visualizing errors and uncertainty. *IEEE Computer Graphics and Applications* 23 (Sept. 2003), 6–10.

[KMH01] KOSARA R., MIKSCH S., HAUSER H.: Semantic depth of field. In *Proc. of IEEE Symp. on Information Visualization* (2001).

[Koe] KOEHN P.: European Parliament Proceedings Parallel Corpus, 1996-2003.

[Koe04] KOEHN P.: *Pharaoh User Manual and Description for Version 1.2*. USC ISI, August 2004.

[KS97] KEMP T., SCHAAF T.: Estimating confidence using word lattices. In *Proc. of Eurospeech* (1997), pp. 827–830.

[LS06] LIU P., SOONG F. K.: Word graph based speech recognition error correction by handwriting input. In *Proc. of the International Conference on Multimodal Interfaces* (Nov. 2006), pp. 339–346.

[Olt06] OLTEANU M.: Phramer: An open-source statistical phrase-based MT decoder. Software, May 2006.

[RJ99] RHEINGANS P., JOSHI S.: Visualization of molecules with positional uncertainty. In *Proc. of Data Visualization* (1999), Springer-Verlag, pp. 299–306.

[SJ00] SMITH N. A., JAHR M. E.: Cairo: An alignment visualization tool. In *Proc. of the Int. Conf. on Language Resources and Evaluation* (2000), pp. 549–552.

[SLK02] SCOTT S. D., LESH N., KLAU G. W.: Investigating human-computer optimization. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems* (Apr. 2002), pp. 155–162.

[SMI99] STROTHOTTE T., MASUCH M., ISENBERG T.: Visualizing knowledge about virtual reconstructions of ancient architecture. In *Proc. of Computer Graphics International* (1999), pp. 36–43.

[SMW01] SUHM B., MYERS B., WAIBEL A.: Multimodal error correction for speech user interfaces. *ACL Transactions on Computer-Human Interaction* 8, 1 (2001), 60–98.

[Sto02] STOLCKE A.: SRILM – an extensible language modeling toolkit. In *Proc. of the Int. Conf. on Spoken Language Processing* (2002), vol. 2, pp. 901–904.

[Tuf01] TUFTE E. R.: *The Visual Display of Quantitative Information*, 2nd ed. Graphics Press, Cheshire, USA, 2001.

[VGRH03] VALTCHEV P., GROSSER D., ROUME C., HACENE M. R.: Galicia: an open platform for lattices. In *Contributions to the 11th Conference on Conceptual Structures* (2003), de Moor A., Lex W., Ganter B., (Eds.), Verlag Shaker, pp. 241–254.

[War04] WARE C.: *Information Visualization: Perception for Design*, 2nd ed. Morgan Kaufmann, 2004.

[YI06] YAMASHITA N., ISHIDA T.: Effects of machine translation on collaborative work. In *Proc. of Computer Supported Cooperative Work* (Nov. 2006), pp. 515–523.

[ZC06] ZUK T., CARPENDALE S.: Theoretical analysis of uncertainty visualizations. In *Proc. of SPIE-IS&T Electronic Imaging* (2006), Erbacher R. F., Roberts J. C., Gröhn M. T., Börner K., (Eds.), vol. 6060, 606007.