Technology for Writing

& what we can learn from theories of cognition and creativity

Candidacy Exam of Katy Ilonka Gero
Department of Computer Science at Columbia University
Summer of COVID-19, 2020

Welcome to a virtual candidacy exam.

Please bring your own snacks.

What is the state of HCI work on writing support?

Writing is not simply transcribing; it is an act of meaning-making and creativity.

So we must also consider work in psychology on writing and creativity.

Outline

1. Theoretical underpinnings.

Models of writing, creativity, and creativity support tools.

2. The “process” view of writing support.

Review of system papers from a process model perspective.

3. The “evaluation” view of writing support.

Review of system papers from an evaluation perspective.

4. Where do we go from here?

Reflection on where there are opportunities for new work.

Outline

1. Theoretical underpinnings.

Models of writing, creativity, and creativity support tools.

2. The “process” view of writing support.

Review of system papers from a process model perspective.

3. The “evaluation” view of writing support.

Review of system papers from an evaluation perspective.

4. Where do we go from here?

Reflection on where there are opportunities for new work.

Writing and creativity will neatly align.

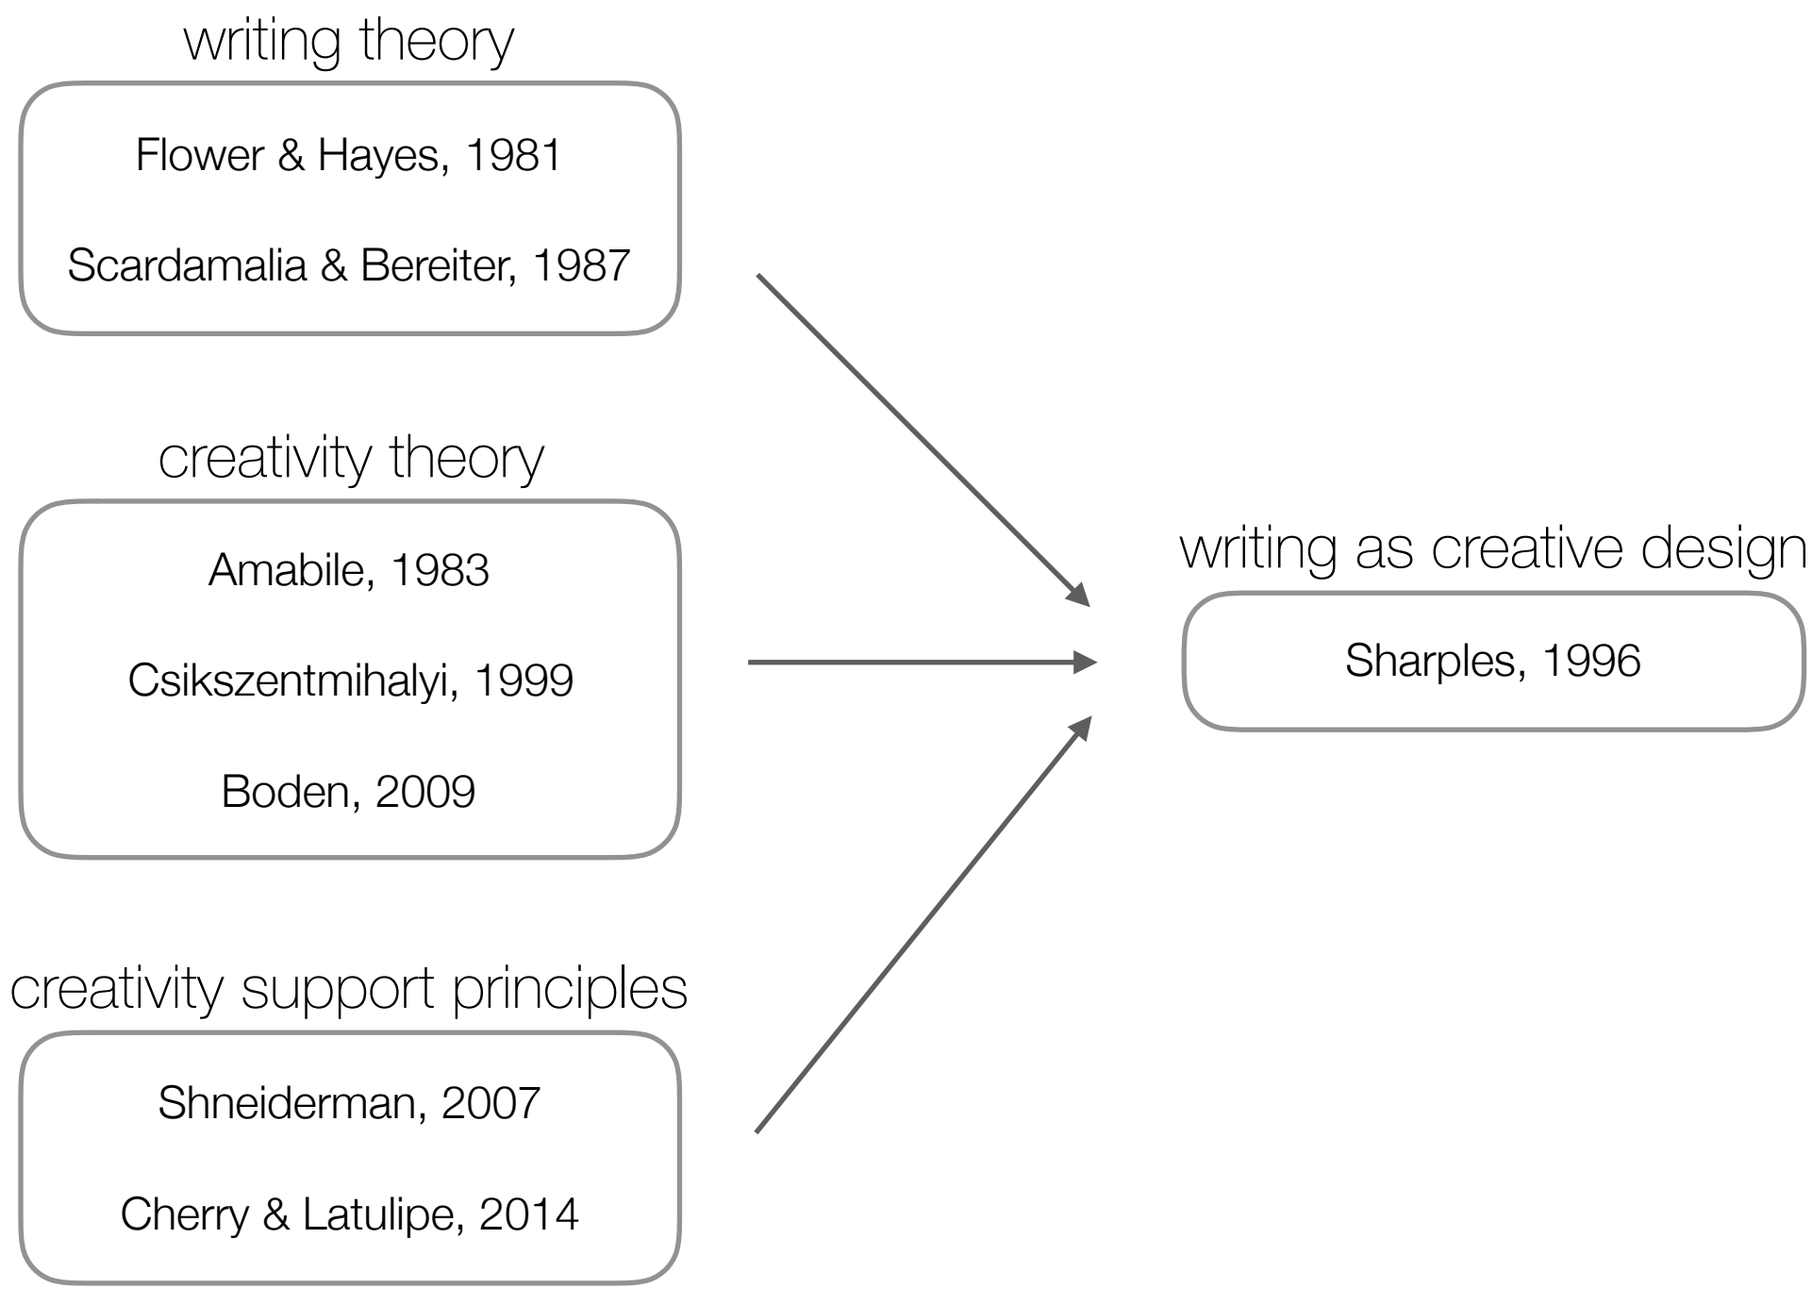

Writing as a set of embedded, cognitive processes.

writing theory

Flower & Hayes, 1981

Scardamalia & Bereiter, 1987

Writing as a set of embedded, cognitive processes.

Flower & Hayes, 1981

A process model of writing.

- writing is a series of nonlinear, embedded cognitive processes
- goals change as writers learn from what they have written

Writing as a set of embedded, cognitive processes.

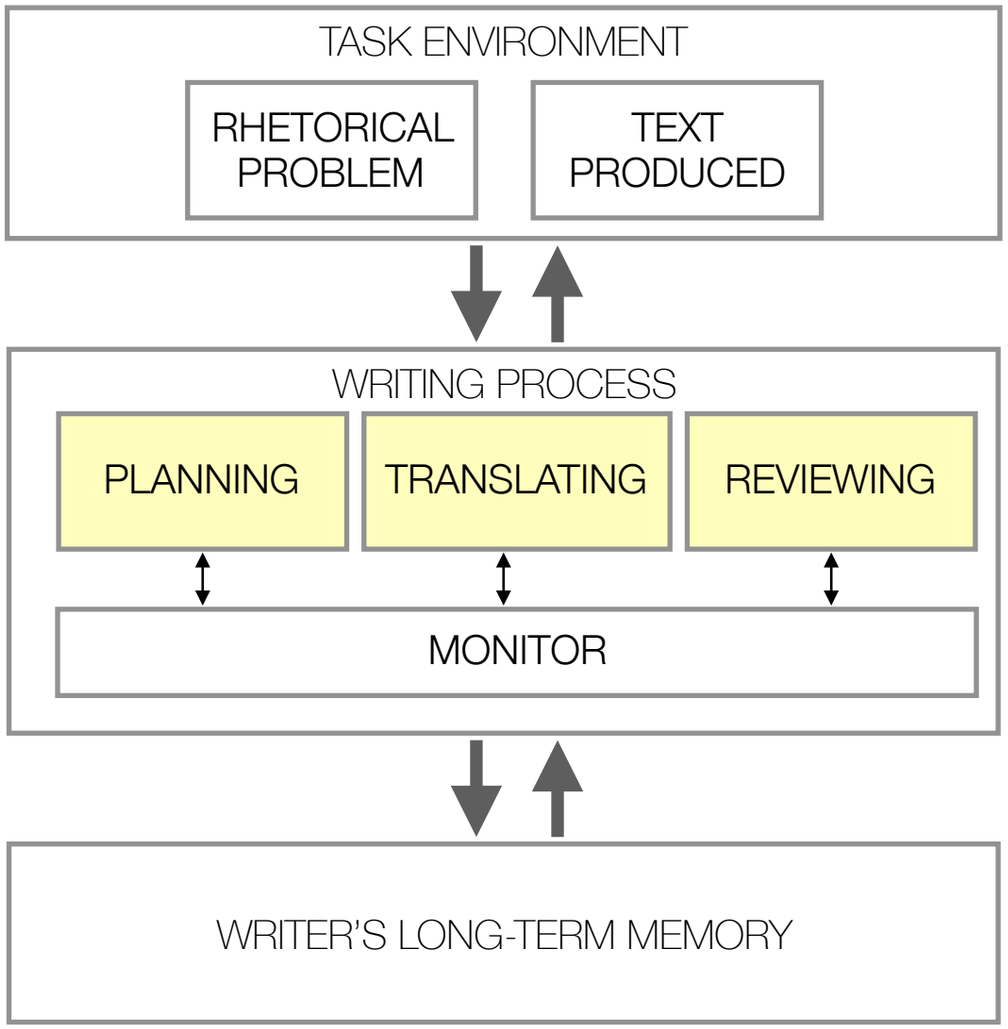

Flower & Hayes, 1981

A process model of writing.

- writing is a series of nonlinear, embedded cognitive processes
- goals change as writers learn from what they have written

Writing as a set of embedded, cognitive processes.

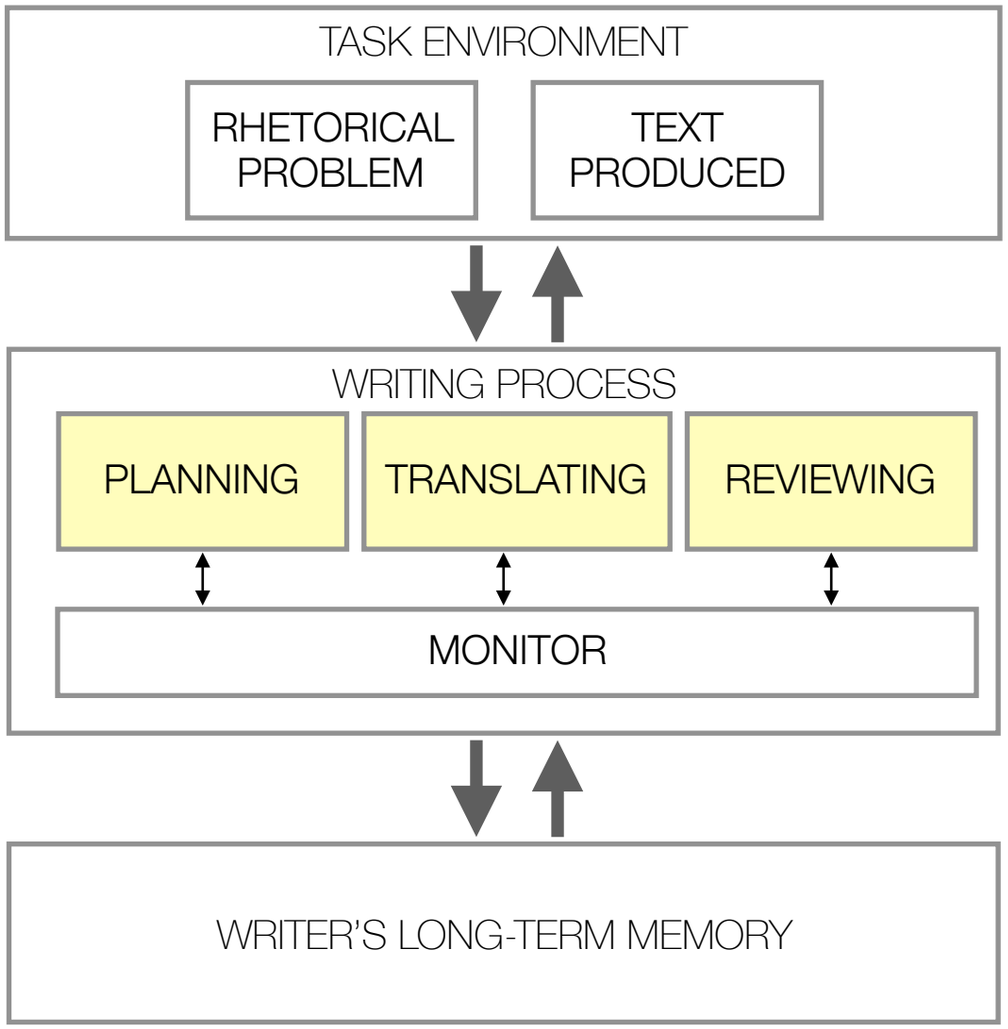

Flower & Hayes, 1981

A process model of writing.

- writing is a series of nonlinear, embedded cognitive processes
- goals change as writers learn from what they have written

Scardamalia & Bereiter, 1987

How do the processes of immature & mature writers differ?

- immature writers use knowledge-telling: topic and genre associations allow for automatic coherence
- mature writers use knowledge-transforming: interaction between content and rhetorical thinking results in new knowledge for the writer

Writing is an act of meaning-making.

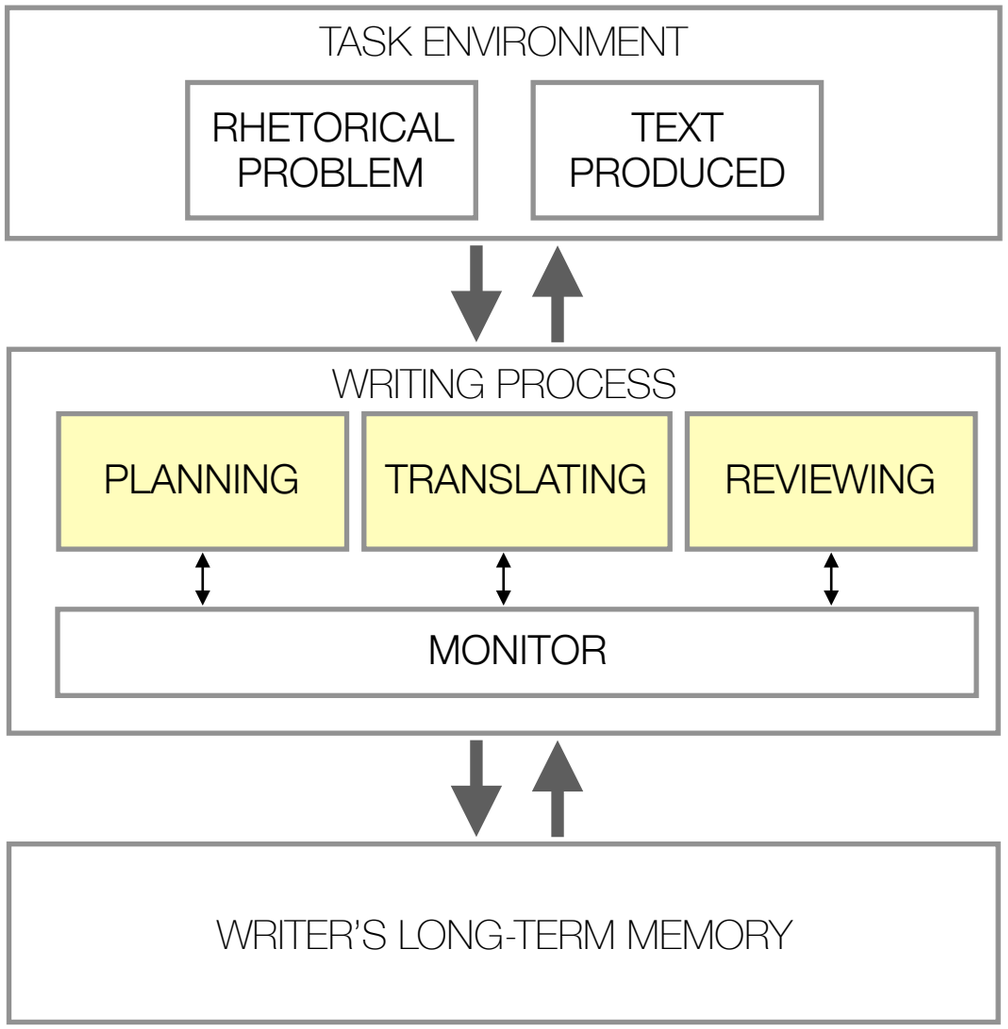

Flower & Hayes, 1981

A process model of writing.

- writing is a series of nonlinear, embedded cognitive processes
- goals change as writers learn from what they have written

Scardamalia & Bereiter, 1987

How do the processes of immature & mature writers differ?

- immature writers use knowledge-telling: topic and genre associations allow for automatic coherence
- mature writers use knowledge-transforming: interaction between content and rhetorical thinking results in new knowledge for the writer

Writing is an act of meaning-making.

writing theory

Flower & Hayes, 1981

Scardamalia & Bereiter, 1987

Technological support for writing must grapple with its complex interaction with thought, and the varied processes involved.

Creativity as a set of embedded, cognitive processes.

writing theory

Flower & Hayes, 1981

Scardamalia & Bereiter, 1987

creativity theory

Amabile, 1983

Csikszentmihalyi, 1999

Boden, 2009

Creativity as a set of embedded, cognitive processes.

writing theory

Flower & Hayes, 1981

Scardamalia & Bereiter, 1987

creativity theory

Amabile, 1983

Csikszentmihalyi, 1999

Boden, 2009

Amabile, 1983

A componential & process model of creativity.

- components: domain-relevant skills, creativity-relevant skills, & task motivation
- processes: task representation, preparation, generation, and validation; can be embedded hierarchically

Creativity as a culturally-embedded process.

writing theory

Flower & Hayes, 1981

Scardamalia & Bereiter, 1987

creativity theory

Amabile, 1983

Csikszentmihalyi, 1999

Boden, 2009

Amabile, 1983

A componential & process model of creativity.

- components: domain-relevant skills, creativity-relevant skills, & task motivation
- processes: task representation, preparation, generation, and validation; can be embedded hierarchically

Csikszentmihalyi, 1999

A systems model of creativity involves the cultural.

- culture (domain); creativity takes place within domains: systems of notation
- social (field); creativity must be socially supported and validated by domain experts
- personal background (individual); creativity may be correlated with certain traits

- to foster creativity, focus on communities rather than individuals

Creativity as a set of computational processes.

writing theory

Flower & Hayes, 1981

Scardamalia & Bereiter, 1987

creativity theory

Amabile, 1983

Csikszentmihalyi, 1999

Boden, 2009

Boden, 2009

A computational model of creativity.

- p-creative versus h-creative
- proposes three “ways” to creativity:
  - <u>combinational</u> produces unfamiliar combinations of familiar ideas
  - <u>exploratory</u> discovers the potentials and limits of a conceptual space
  - <u>transformational</u> alters defining dimensions of a conceptual space

Creativity is process-driven and culturally-entwined.

writing theory

Flower & Hayes, 1981

Scardamalia & Bereiter, 1987

creativity theory

Amabile, 1983

Csikszentmihalyi, 1999

Boden, 2009

Creativity can also be decomposed into processes, and must also grapple with complex interaction with culture.

Design principles for creativity support tools.

writing theory

Flower & Hayes, 1981

Scardamalia & Bereiter, 1987

creativity theory

Amabile, 1983

Csikszentmihalyi, 1999

Boden, 2009

creativity support principles

Shneiderman, 2007

Cherry & Latulipe, 2014

Design principles for creativity support tools.

Shneiderman, 2007

Design principles for creativity support.

- inspired by work on theories of creativity:
  - enable collaboration
  - support exploratory search
  - low thresholds, high ceilings, and wide walls
  - provide rich history-keeping
  - evaluation moves toward case studies

Design principles for creativity support tools.

Shneiderman Principles

COLLABORATION

EXPLORATORY SEARCH

WIDE WALLS

LOW THRESHOLDS,
HIGH CEILINGS

HISTORY-KEEPING

Shneiderman, 2007

Design principles for creativity support.

- inspired by work on theories of creativity:
  - enable collaboration
  - support exploratory search
  - low thresholds, high ceilings, and wide walls
  - provide rich history-keeping
- evaluation moves toward case studies

Design principles for creativity support tools.

Shneiderman Principles

COLLABORATION

EXPLORATORY SEARCH

WIDE WALLS

LOW THRESHOLDS, HIGH CEILINGS

HISTORY-KEEPING

Shneiderman, 2007

Design principles for creativity support.

- inspired by work on theories of creativity:
  - enable collaboration
  - support exploratory search
  - low thresholds, high ceilings, and wide walls
  - provide rich history-keeping
- evaluation moves toward case studies

Cherry & Latulipe, 2014

A creativity support index for evaluation.

- measures and weights six factors:
  - collaboration, enjoyment, exploration, expressiveness, immersion & results worth effort
- an individual score represents the intersection of task, tool, and user

Design principles for creativity support tools.

Shneiderman Principles

- COLLABORATION
- EXPLORATORY SEARCH
- WIDE WALLS
- LOW THRESHOLDS, HIGH CEILINGS
- HISTORY-KEEPING

Creativity Support Index

- COLLABORATION
- ENJOYMENT
- EXPLORATION
- EXPRESSIVENESS
- IMMERSION
- RESULTS WORTH EFFORT

Shneiderman, 2007

Design principles for creativity support.

- inspired by work on theories of creativity:
  - enable collaboration
  - support exploratory search
  - low thresholds, high ceilings, and wide walls
  - provide rich history-keeping
- evaluation moves toward case studies

Cherry & Latulipe, 2014

A creativity support index for evaluation.

- measures and weights six factors:
  - collaboration, enjoyment, exploration, expressiveness, immersion & results worth effort
- an individual score represents the intersection of task, tool, and user

Design principles for creativity support tools.

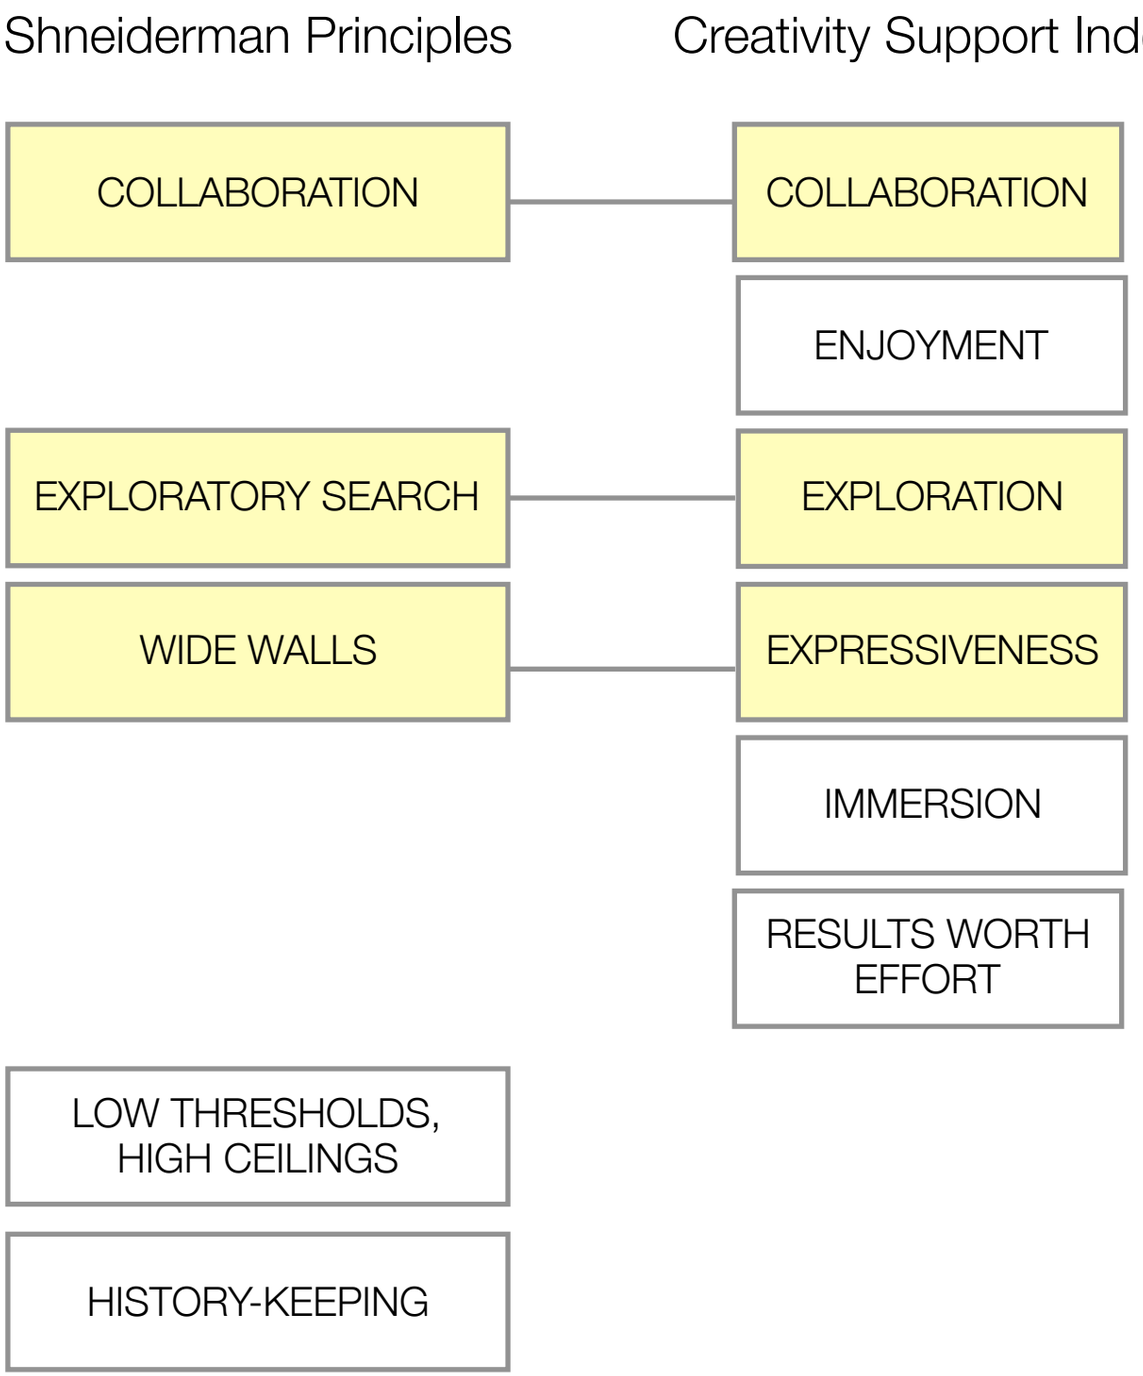

Shneiderman, 2007

Design principles for creativity support.

- inspired by work on theories of creativity:
  - enable collaboration
  - support exploratory search
  - low thresholds, high ceilings, and wide walls
  - provide rich history-keeping
- evaluation moves toward case studies

Cherry & Latulipe, 2014

A creativity support index for evaluation.

- measures and weights six factors:
  - collaboration, enjoyment, exploration, expressiveness, immersion & results worth effort
- an individual score represents the intersection of task, tool, and user

Design principles for creativity support tools.

writing theory

Flower & Hayes, 1981

Scardamalia & Bereiter, 1987

creativity theory

Amabile, 1983

Csikszentmihalyi, 1999

Boden, 2009

creativity support principles

Shneiderman, 2007

Cherry & Latulipe, 2014

Creativity support tools can be evaluated rigorously, either through case studies or more quantitative analysis like the CSI.

Writing as an act of creative design.

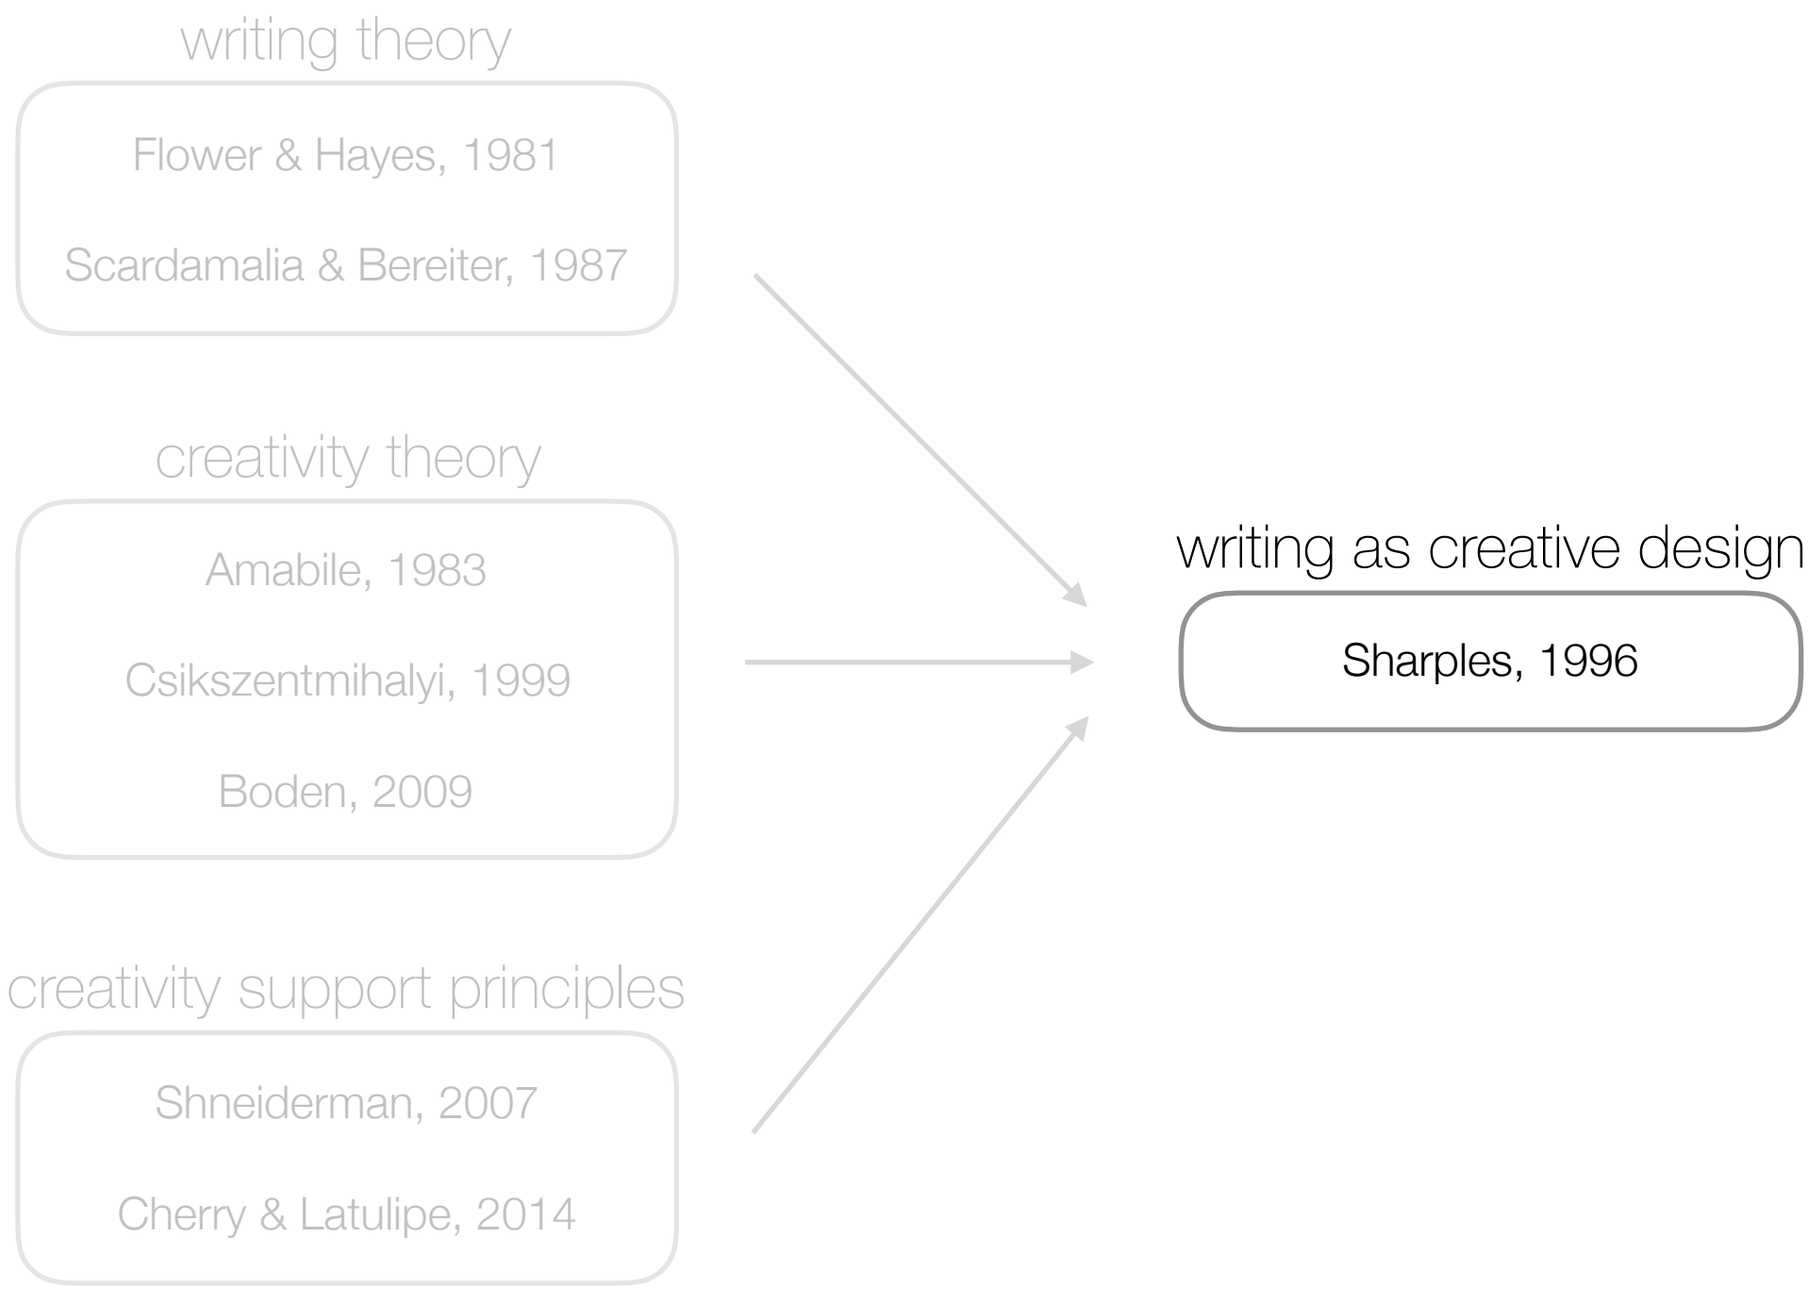

Writing as an act of creative design.

PRIMARY GENERATORS

IMPORTANCE OF ITERATION

TOOLS AND EXTERNAL REPRESENTATIONS

Sharples, 1999

How design theory can inform the writing process.

- primary generators are powerful and early ideas that prompt and guide activity
- iteration through the fusion of analysis (reviewing) and synthesis (translating)
- the writer is a user of tools and creator of cognitive artifacts
- writing as a creative mental process involves movement between:
  - engagement (translating)
  - reflection (reviewing & planning)

Writing as an act of creative design.

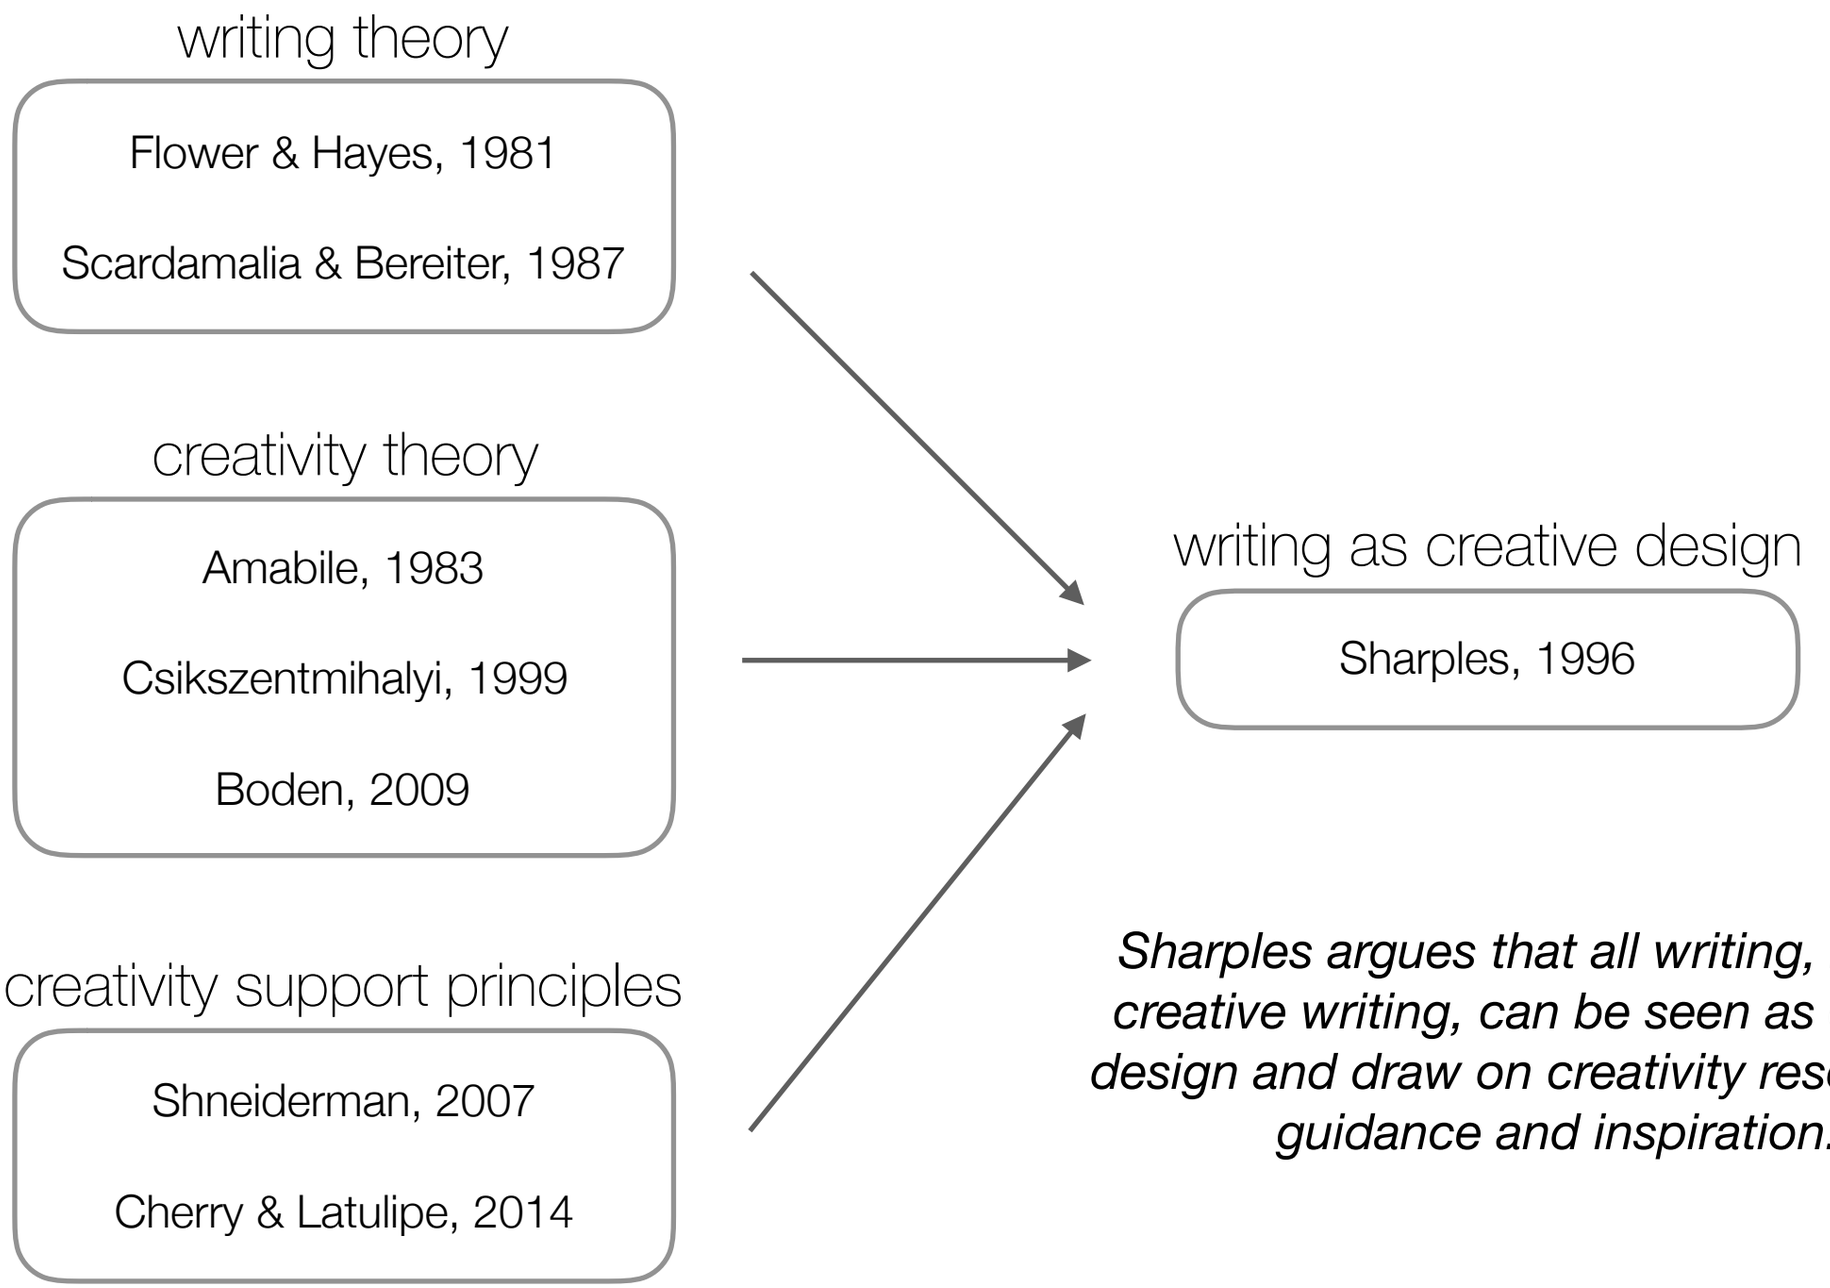

Sharples argues that all writing, not just creative writing, can be seen as creative design and draw on creativity research for guidance and inspiration.

Writing and creativity will neatly align.

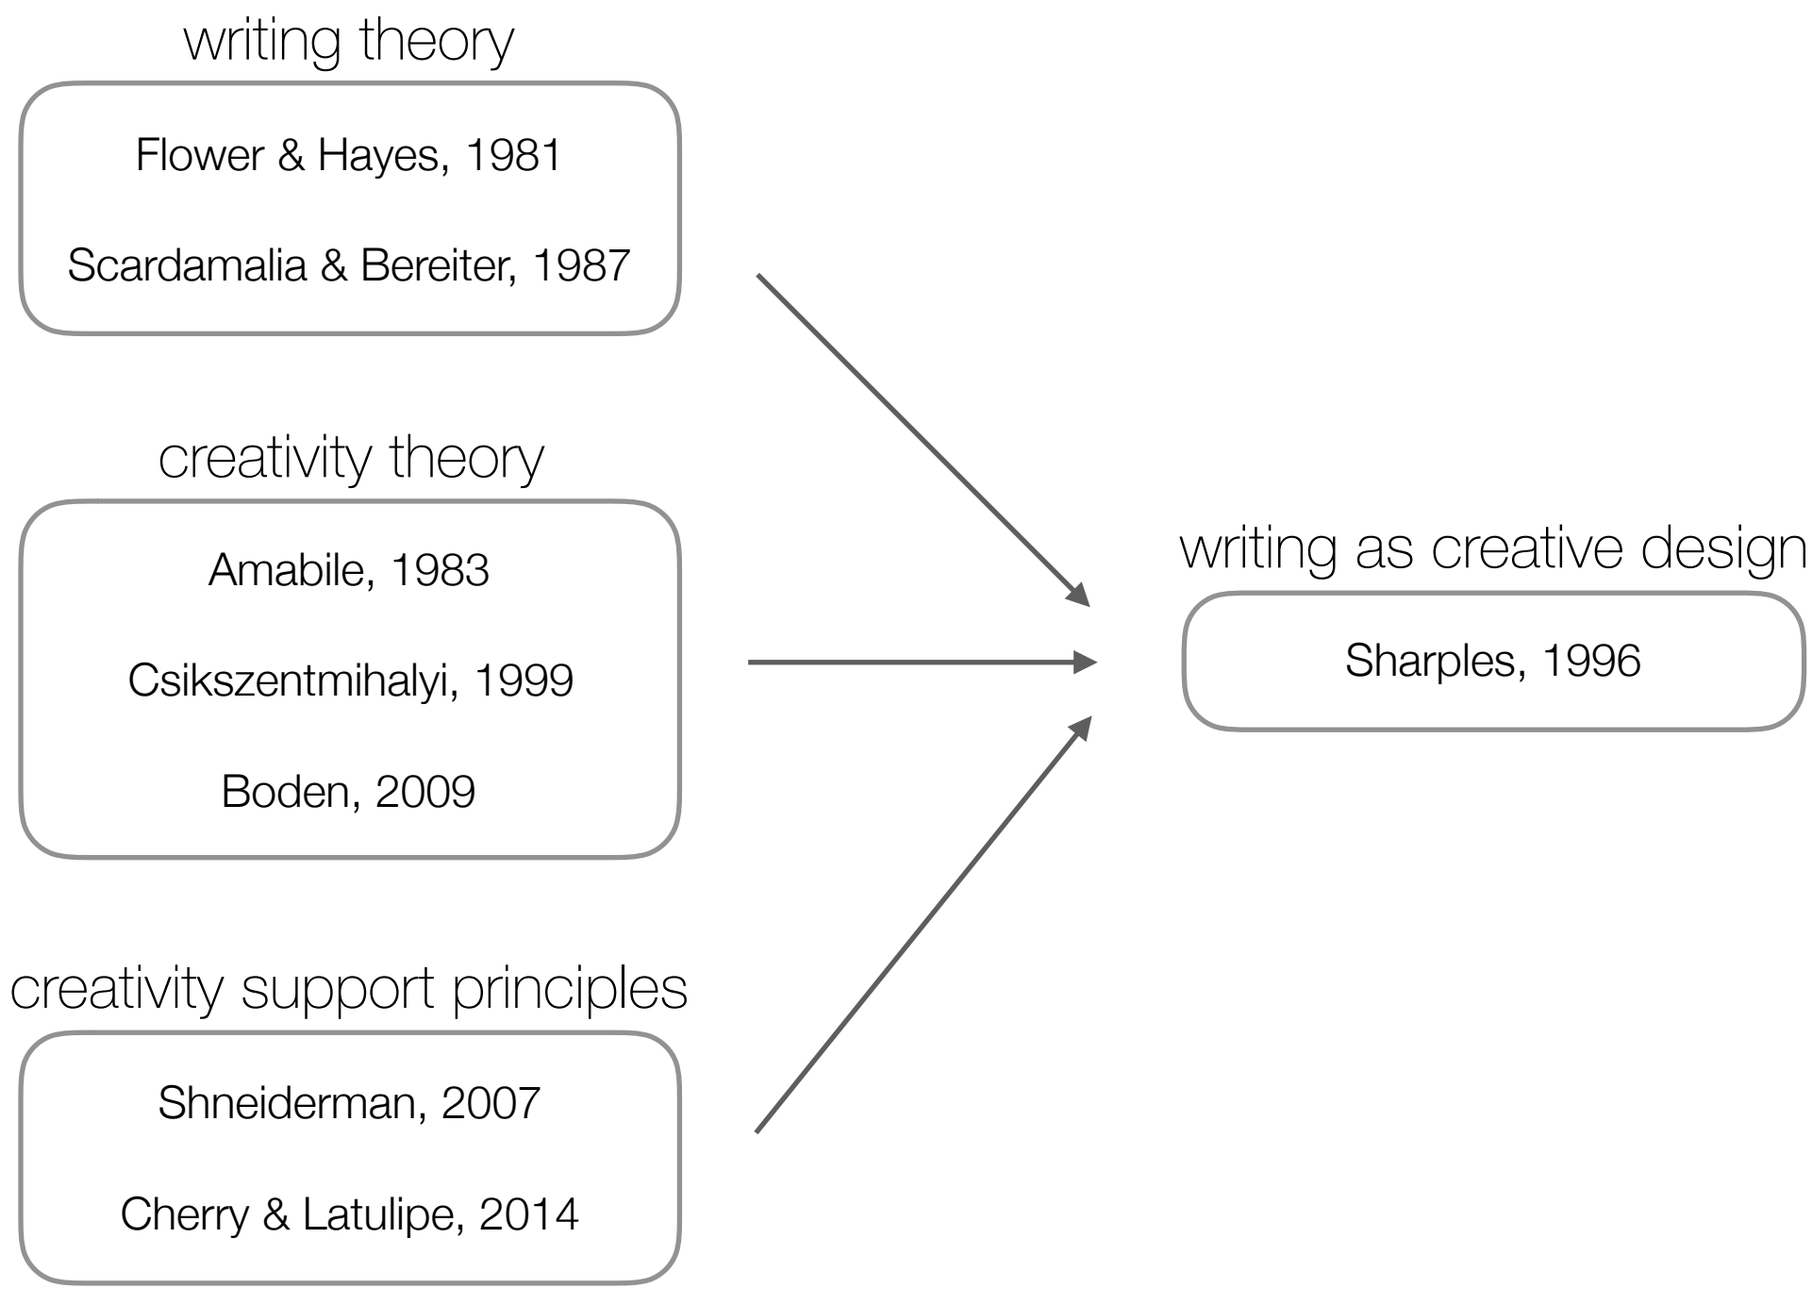

Analogue processes drive writing and creativity.

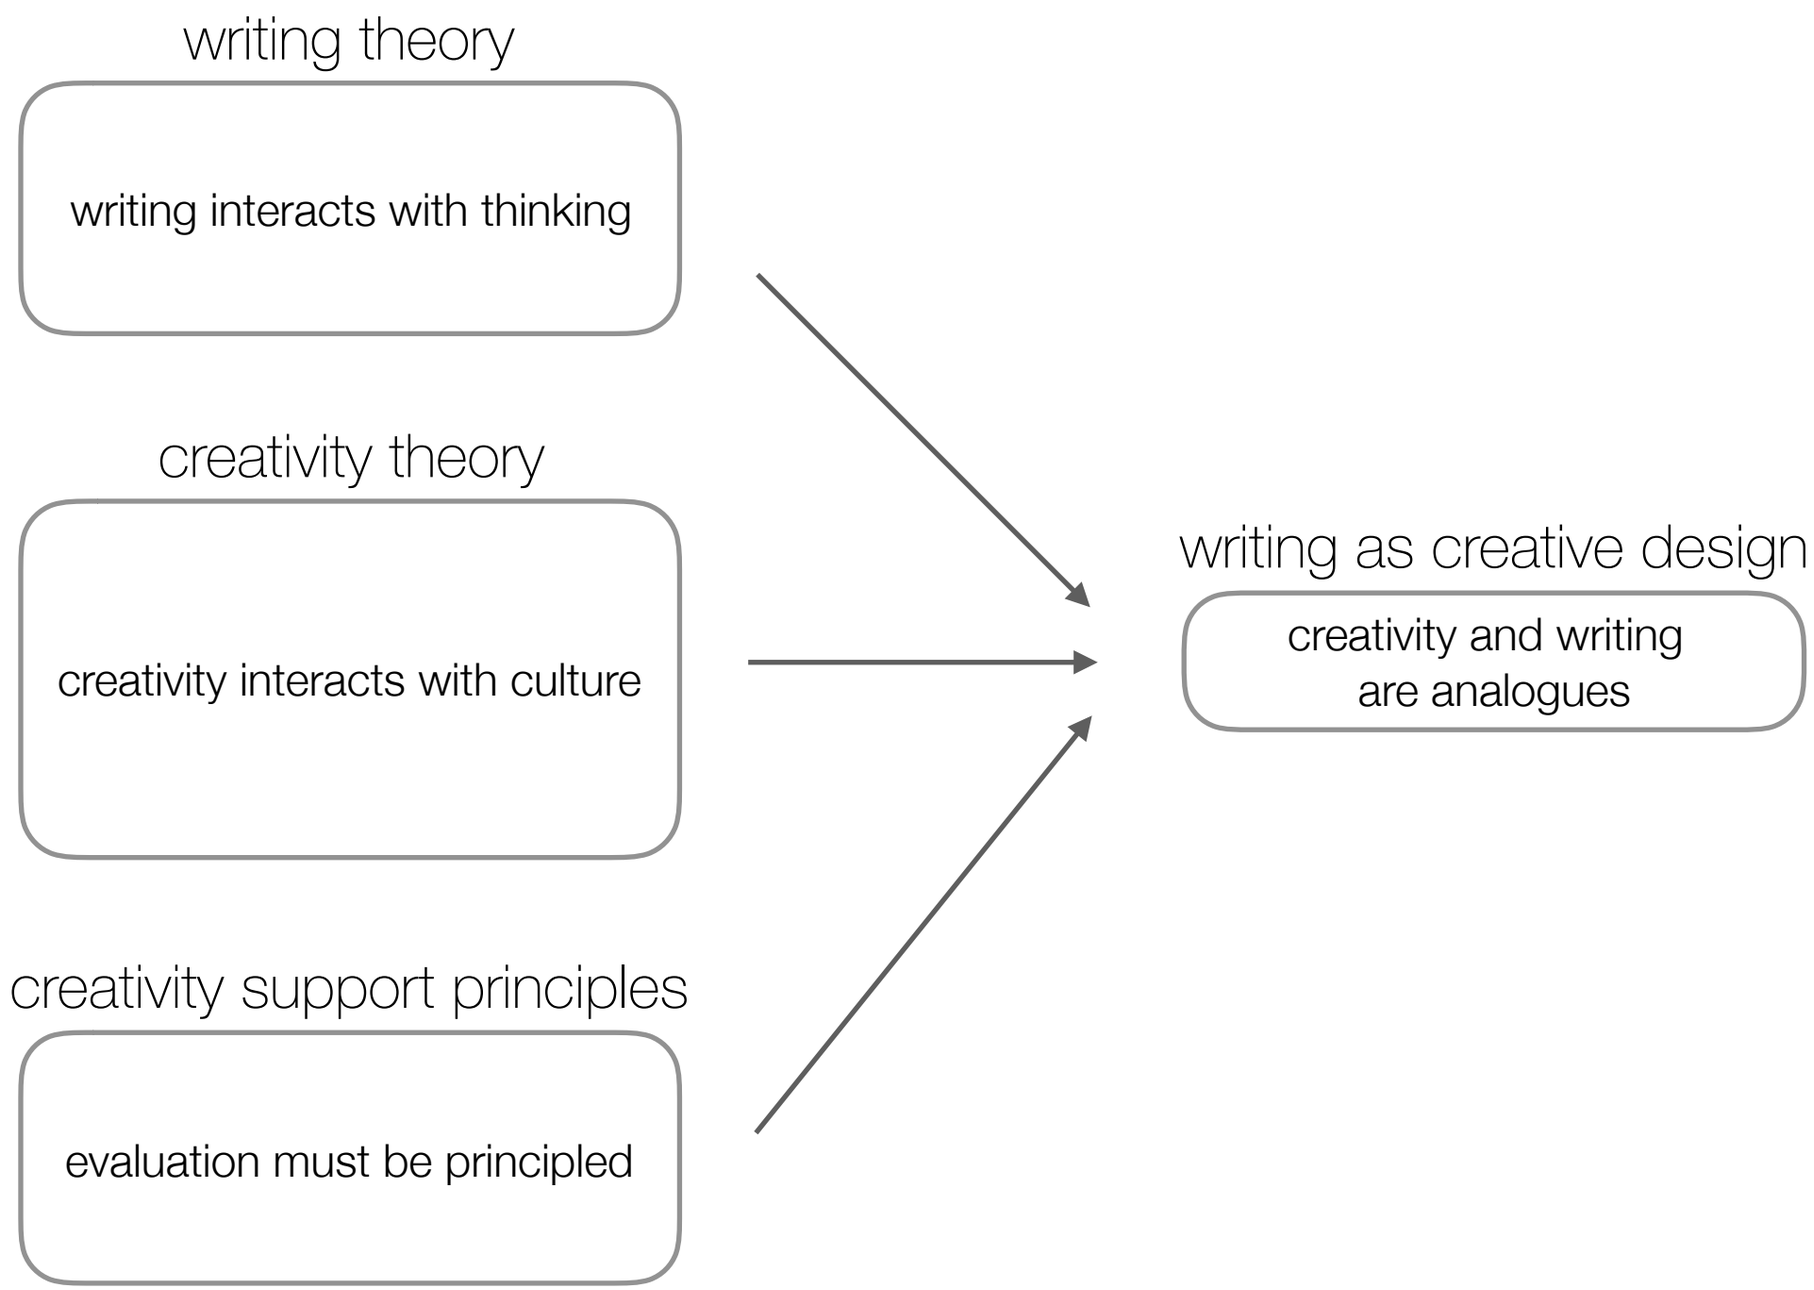

Outline

1. Theoretical underpinnings.

Writing can be viewed as creative design, and we can define nonlinear, yet distinct, cognitive processes: planning, generation, and review.

2. The “process” view of writing support.

Review of system papers from a process model perspective.

3. The “evaluation” view of writing support.

Review of system papers from an evaluation perspective.

4. Where do we go from here?

Reflection on where there are opportunities for new work.

Outline

1. Theoretical underpinnings.

    Writing can be viewed as creative design, and we can define nonlinear, yet distinct, cognitive processes: planning, generation, and review.

2. The “process” view of writing support.

    Review of system papers from a process model perspective.

3. The “evaluation” view of writing support.

    Review of system papers from an evaluation perspective.

4. Where do we go from here?

    Reflection on where there are opportunities for new work.

Writing and creativity as analogues.

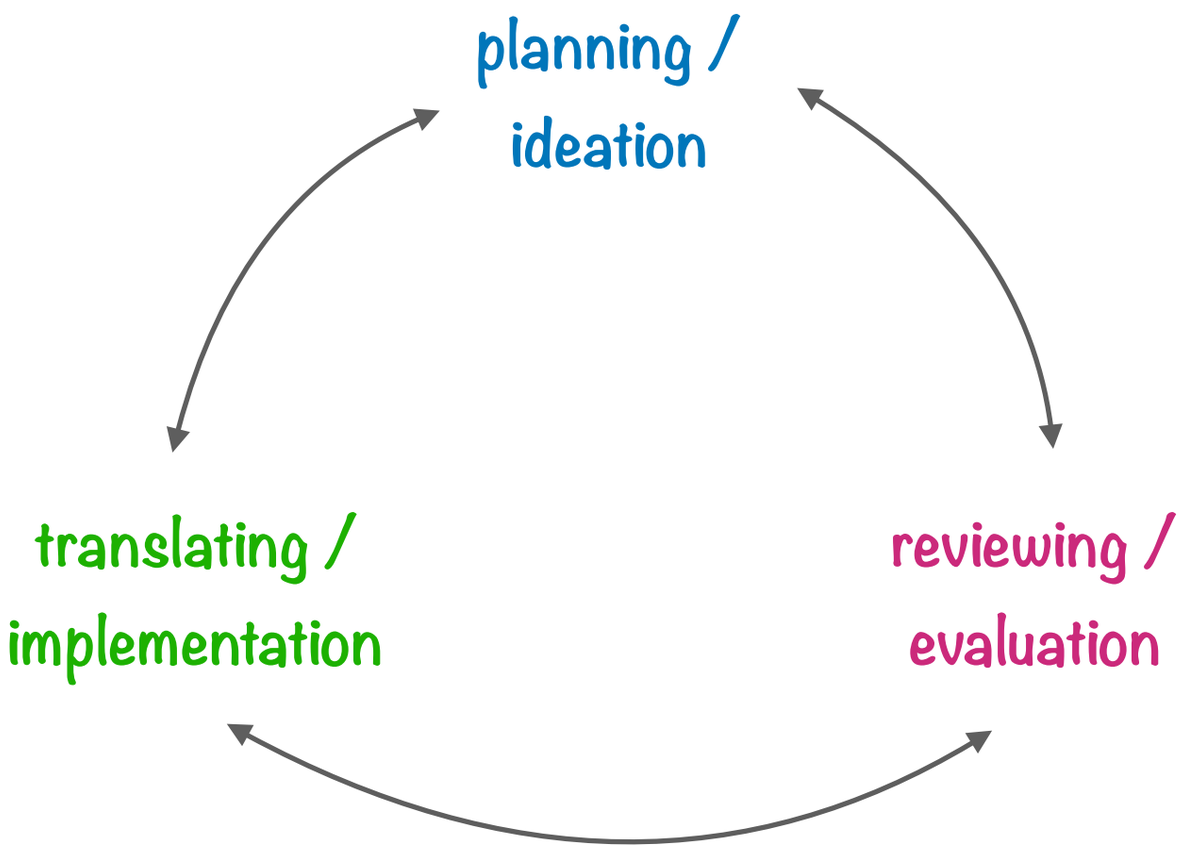

- writing theory points towards a 3-part process model
  - planning, translating, & reviewing
- creativity theory also points to towards a 3-part process model
  - ideation, implementation, & evaluation
- these models neatly align
- they allow us to move beyond what a system *does* to what it *supports*
- some systems support moving between stages

Writing and creativity as analogues.

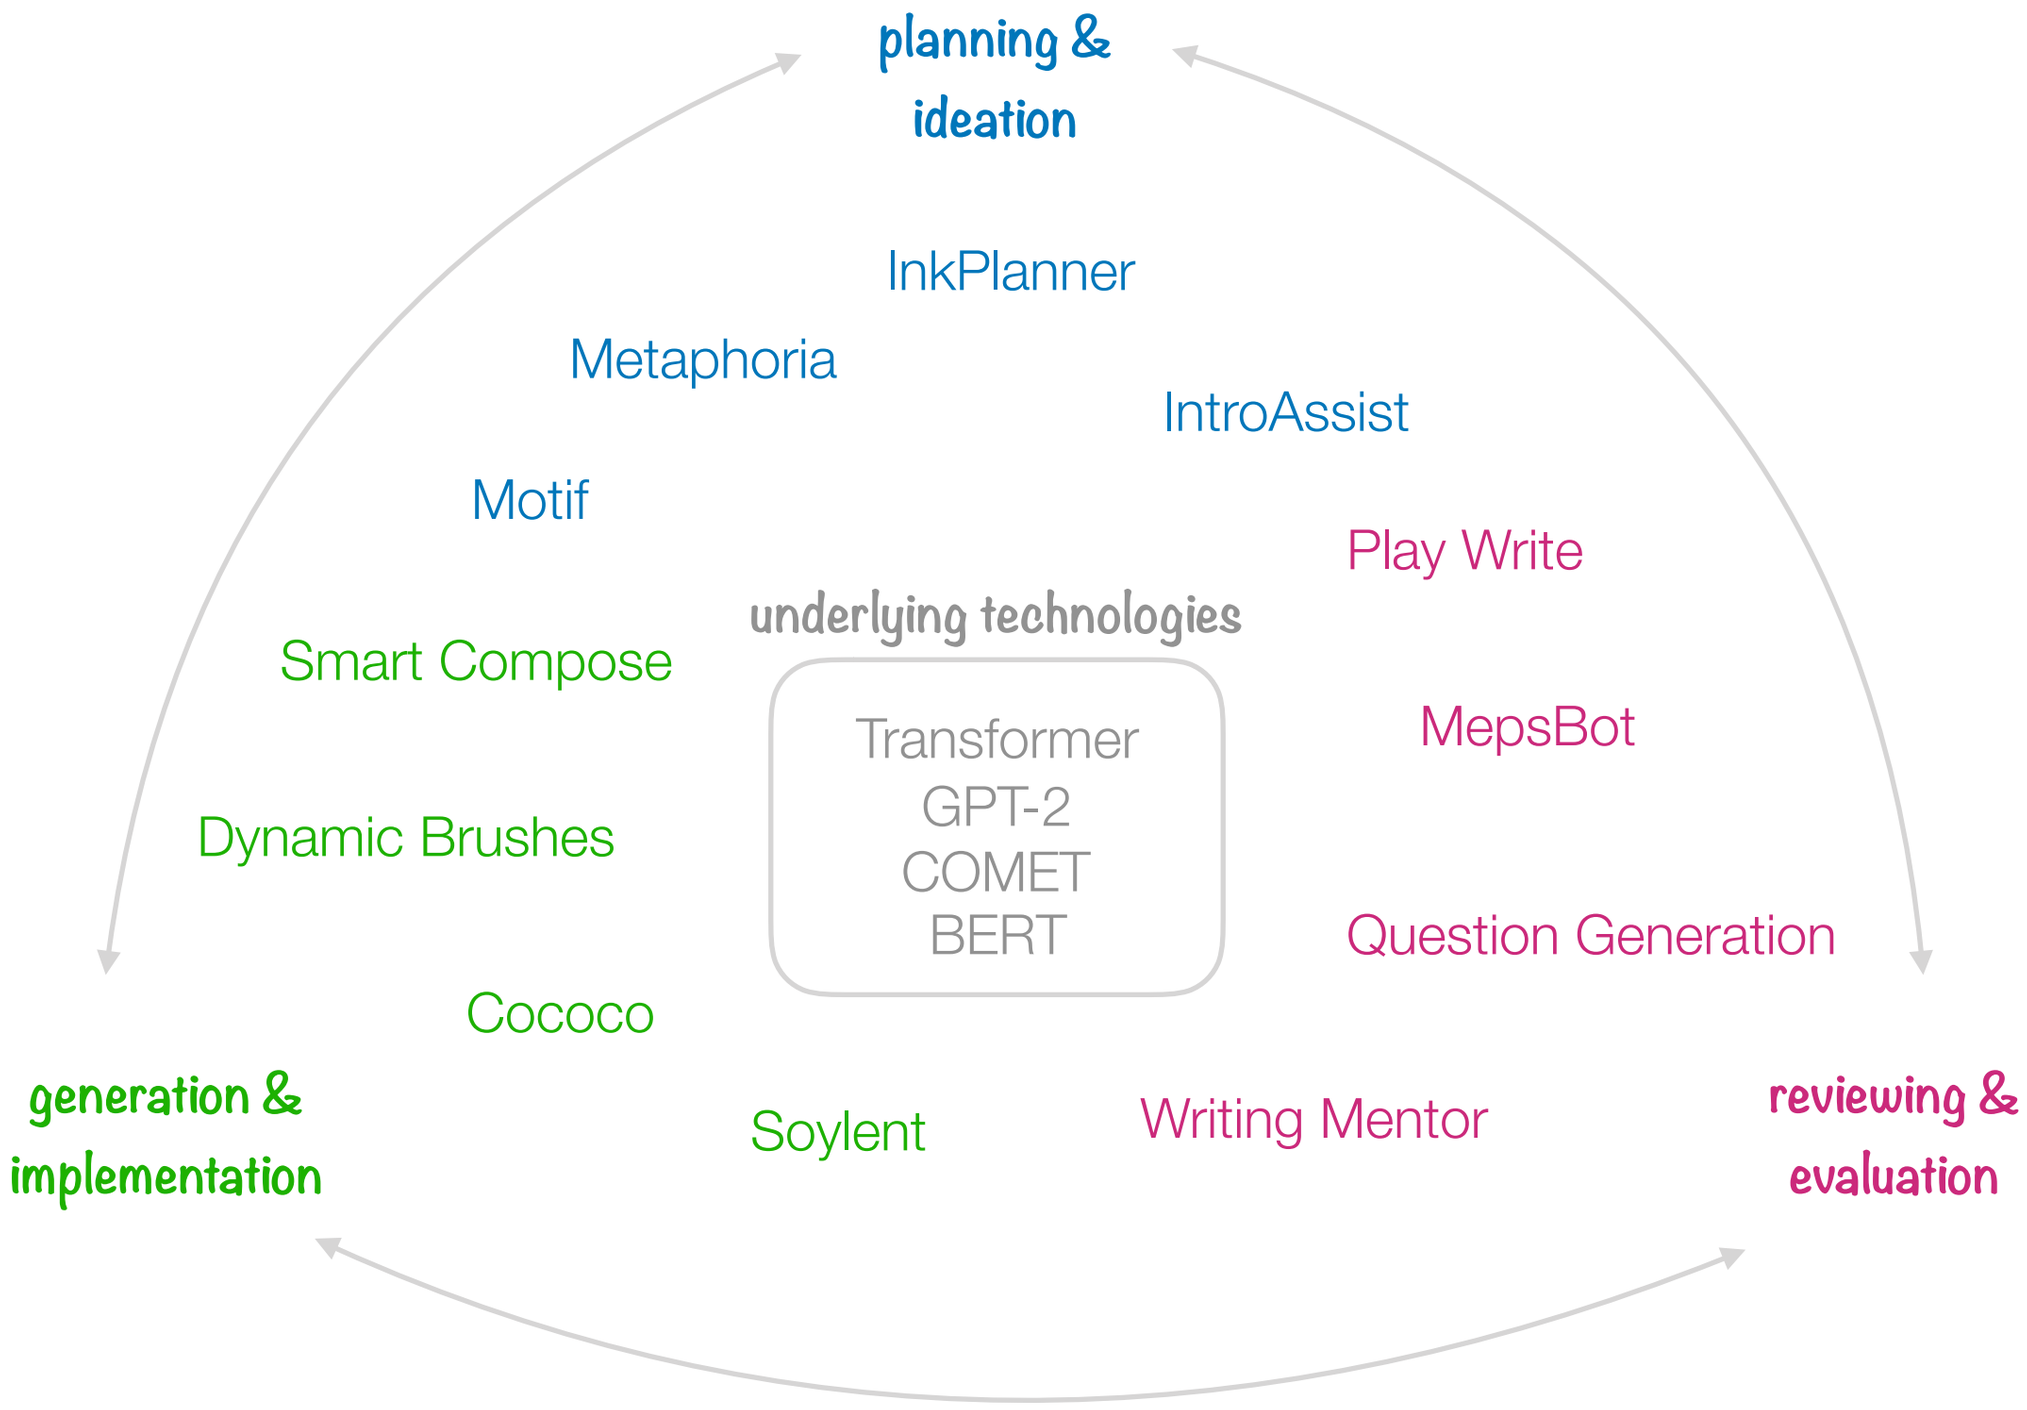

Most system papers are about writing…

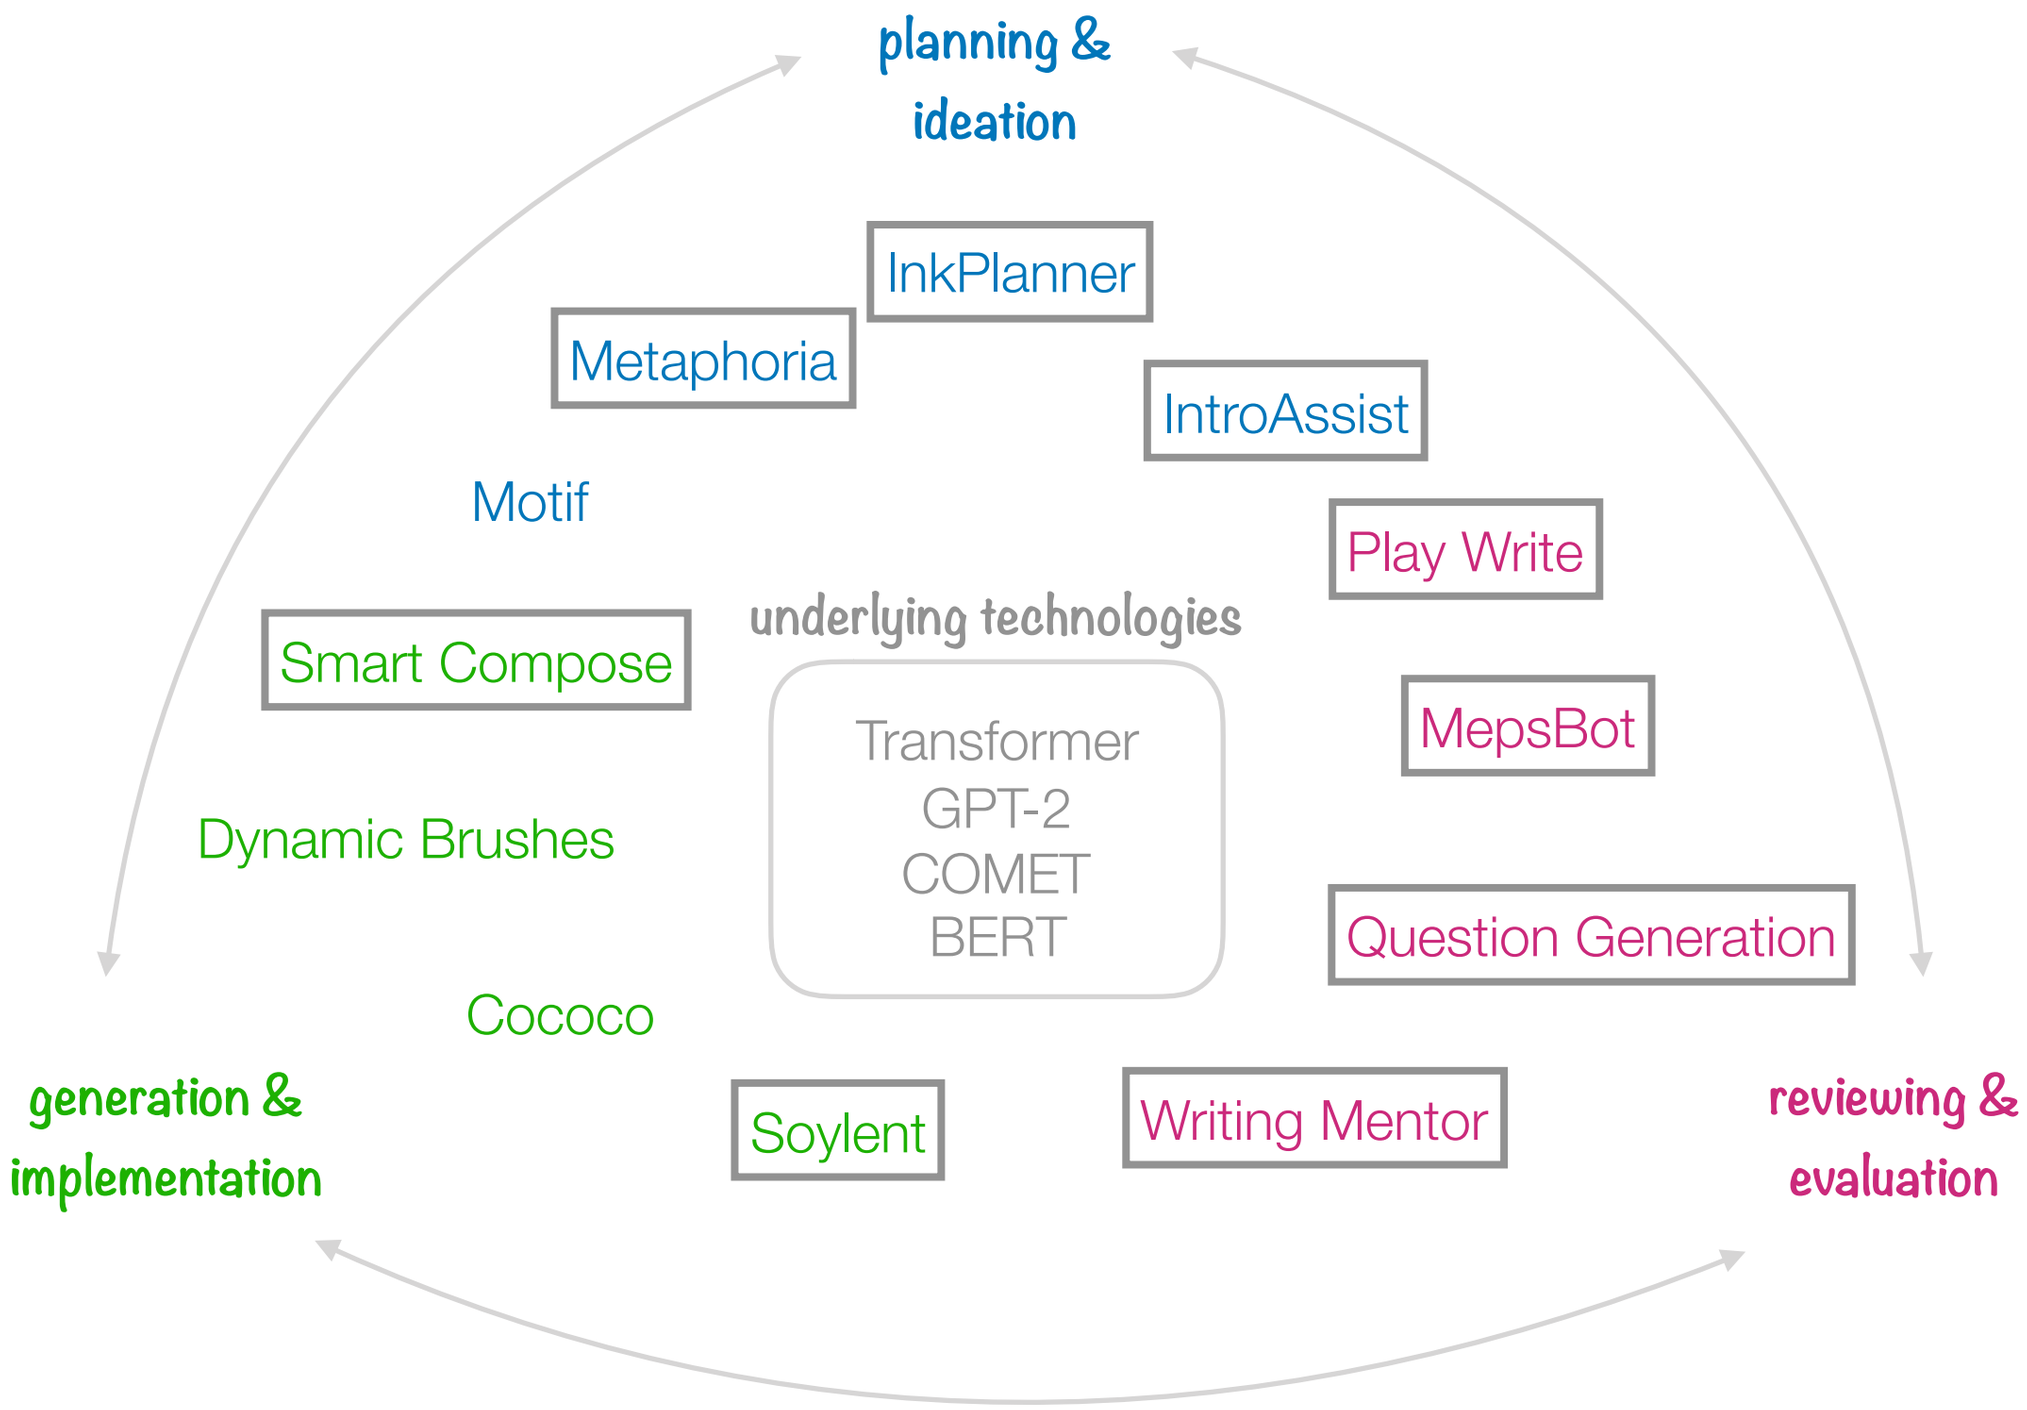

… and systems in other mediums can pave the way.

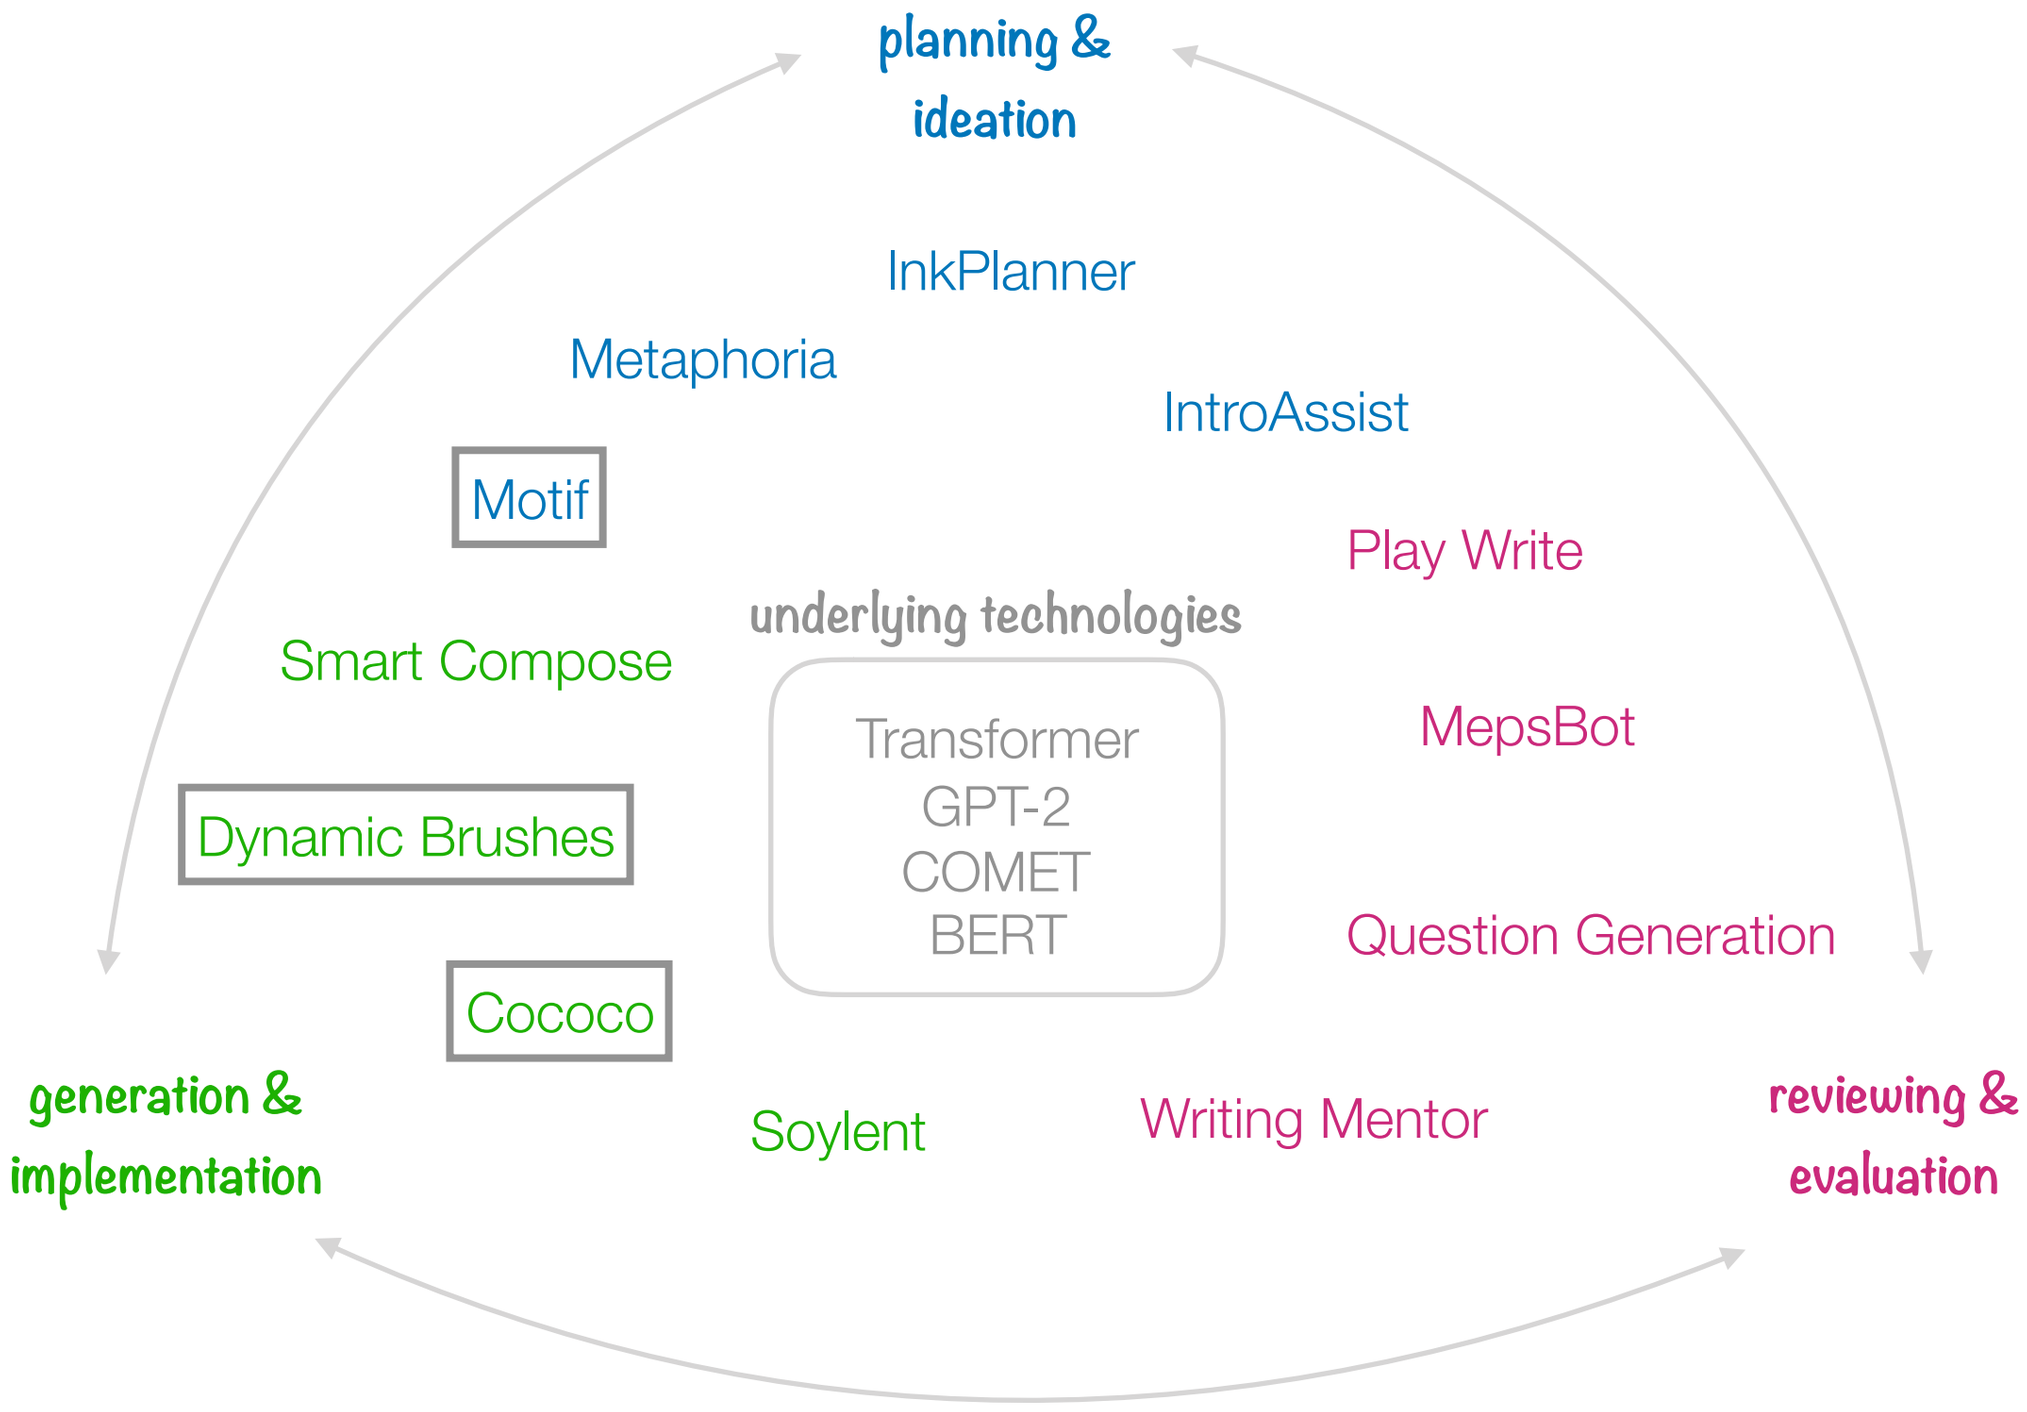

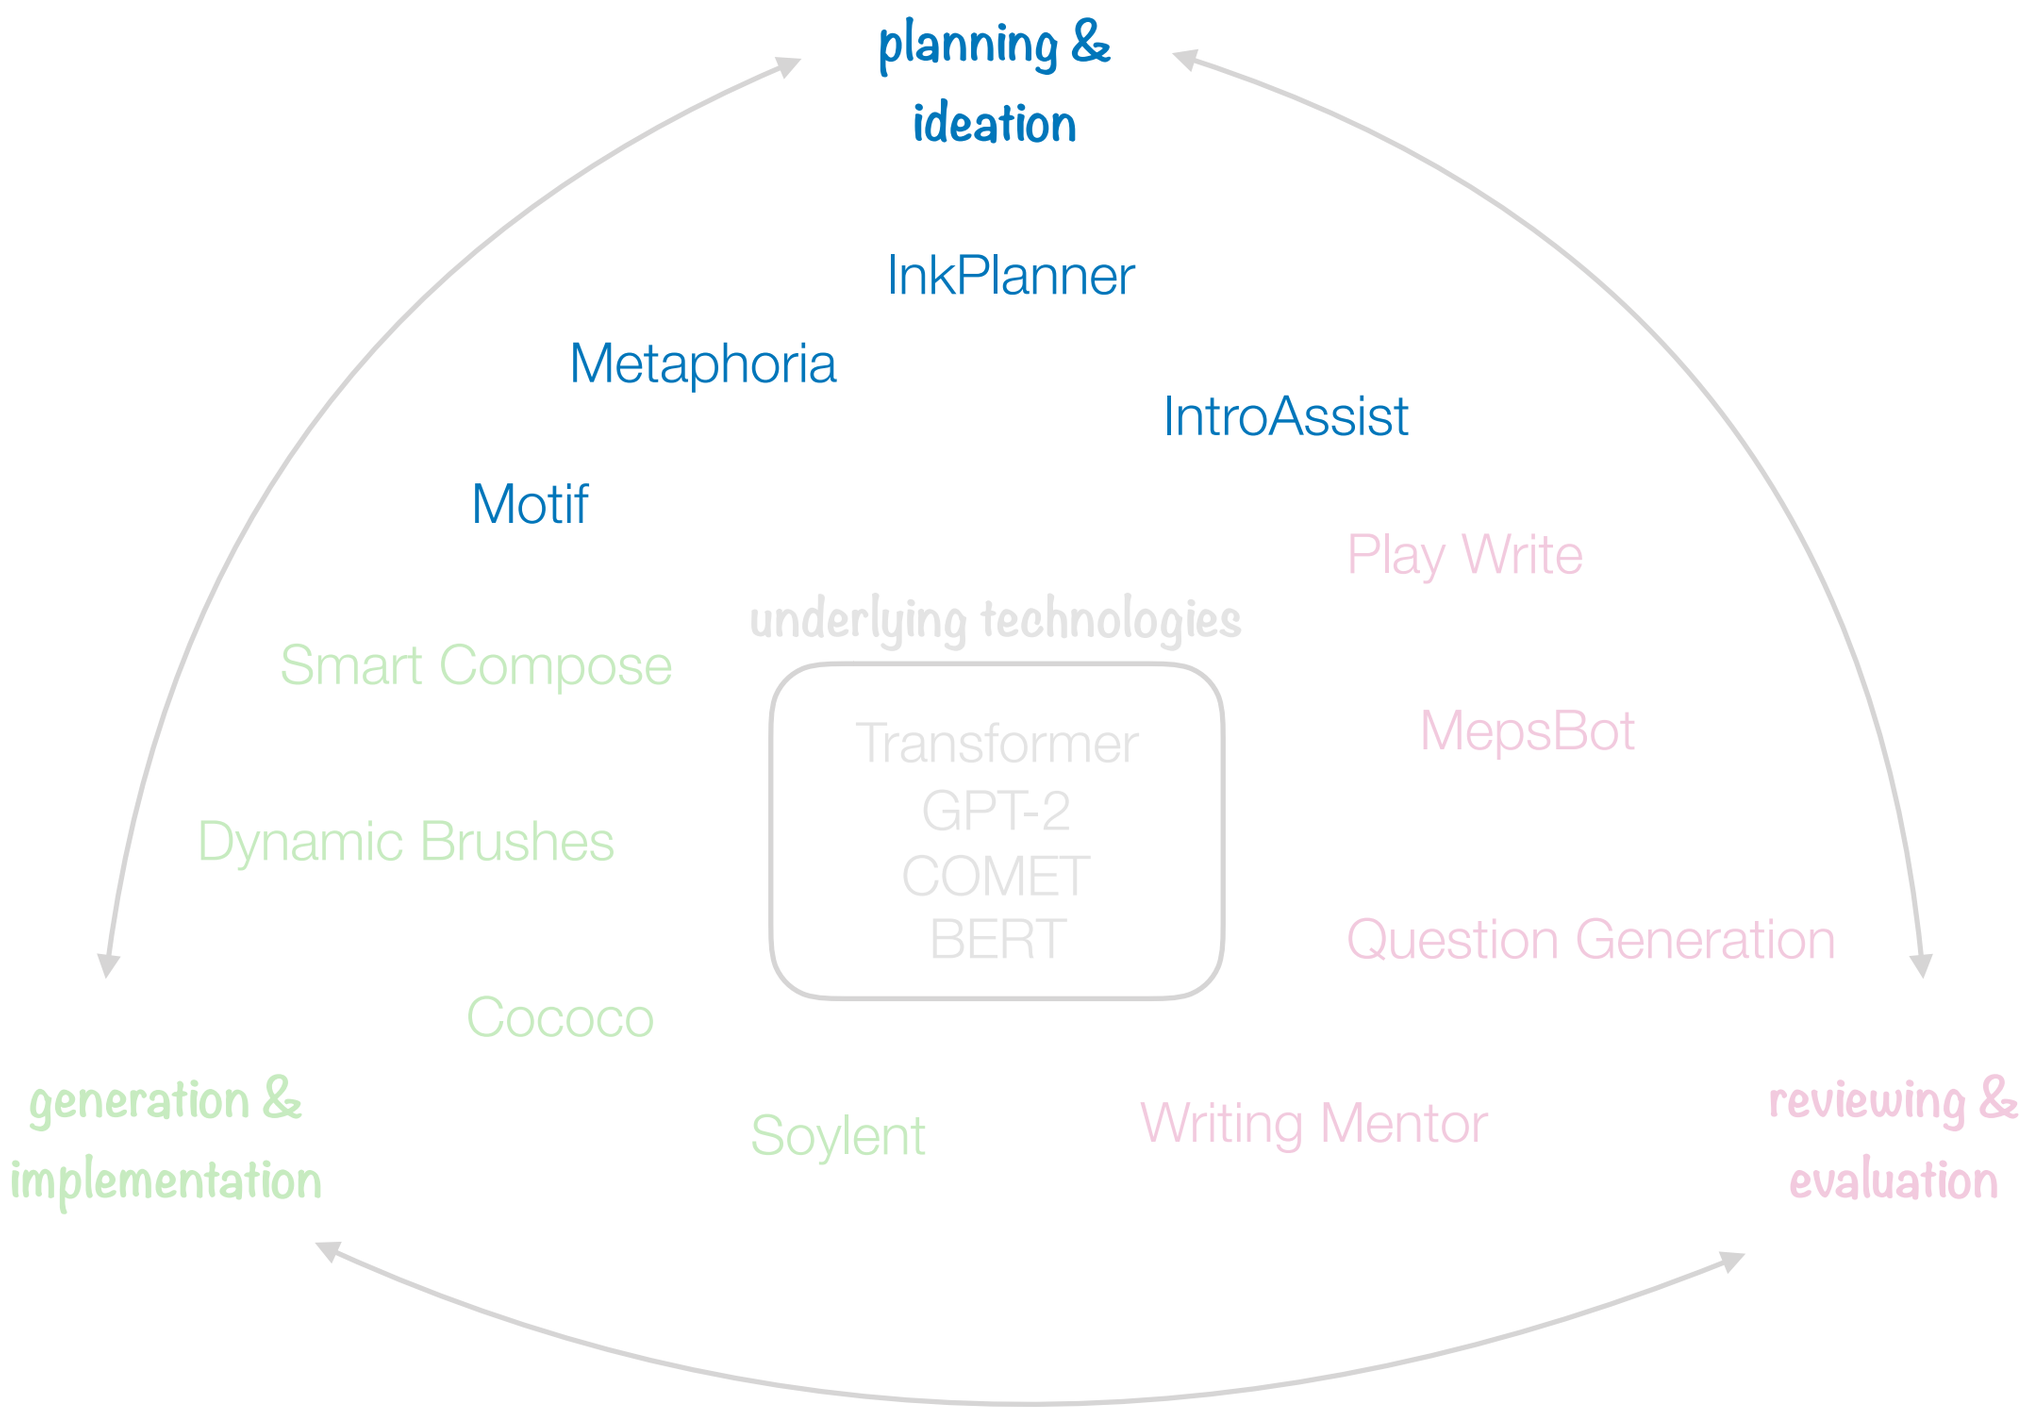
planning & ideation
InkPlanner
Metaphoria
IntroAssist
Motif
Play Write
underlying technologies
Smart Compose
Transformer
GPT-2
COMET
BERT
MepsBot
Dynamic Brushes
Question Generation
Cococo
generation & implementation
Soylent
Writing Mentor
reviewing & evaluation

Support for process versus for semantics.

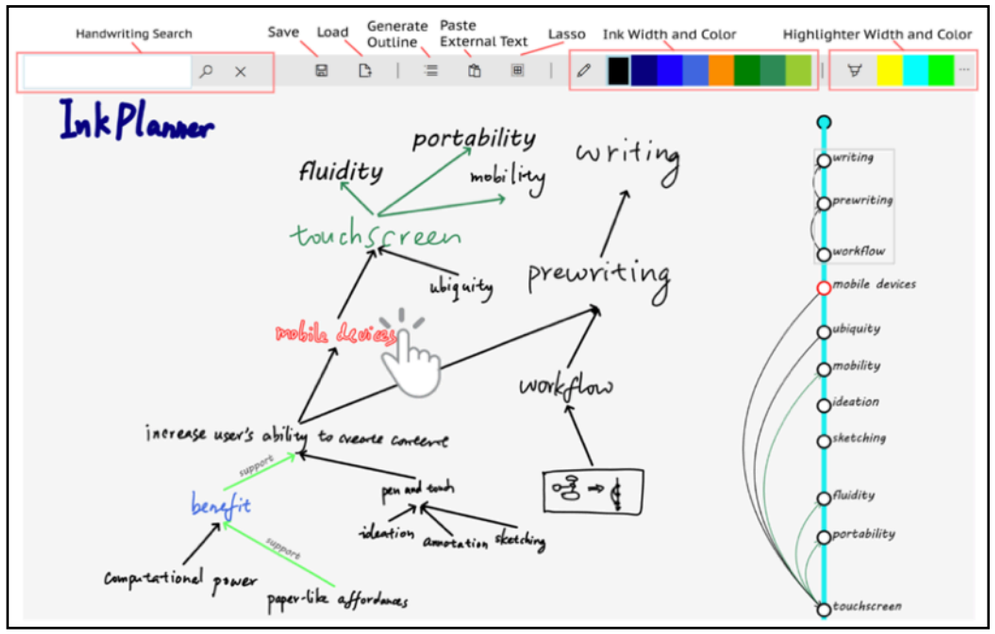

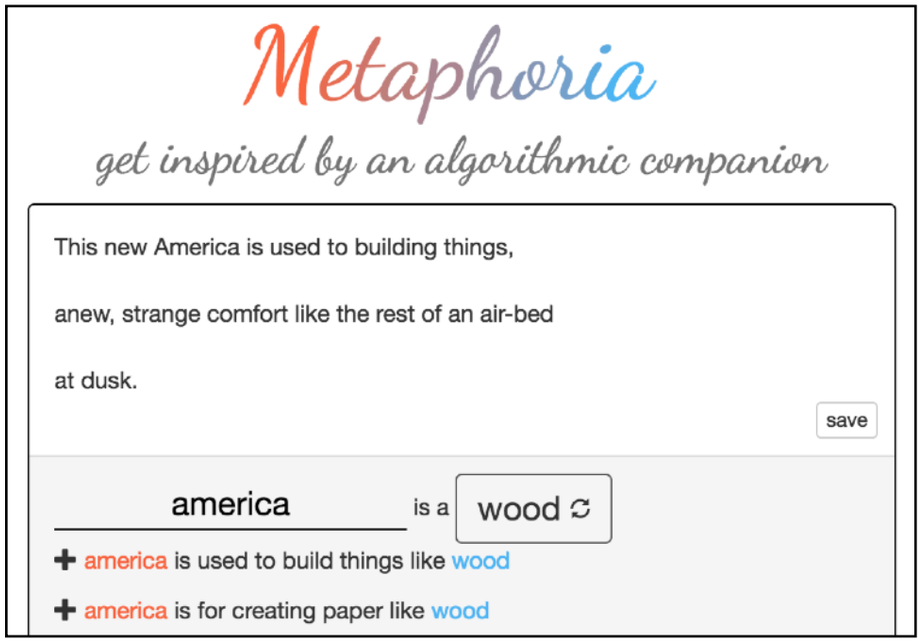

InkPlanner

- supports planning with better diagramming tools
- writer jots down unstructured text on tablet interface; tablet supports mind-mapping and outlining (as well as brainstorming)
- system mainly augments *process*; thus the evaluation focuses on usability

Metaphoria

- supports ideation with suggested metaphorical connections
- writer inputs seed word and requests help with button outside of text-area; tool suggests 10 potential metaphors
- system mainly augments *semantics*; thus the evaluation focuses on quality

Lu, et al. “Inkplanner: Supporting prewriting via intelligent visual diagramming.” TVCG 2018.
Gero & Chilton. “Metaphora: An algorithmic companion for metaphor creation.” CHI 2019.

Grammars & templates for structured forms.

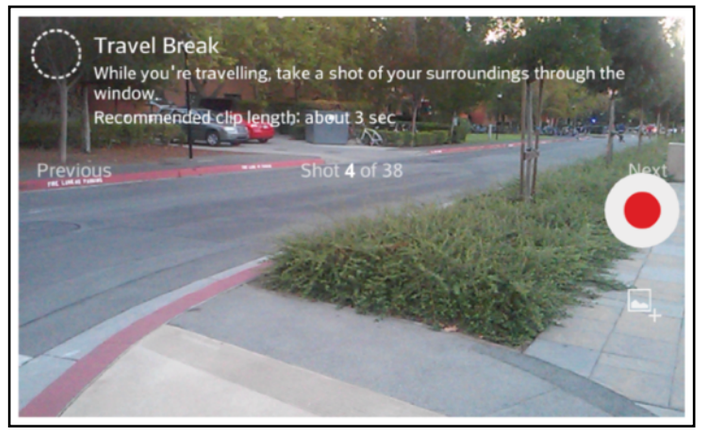

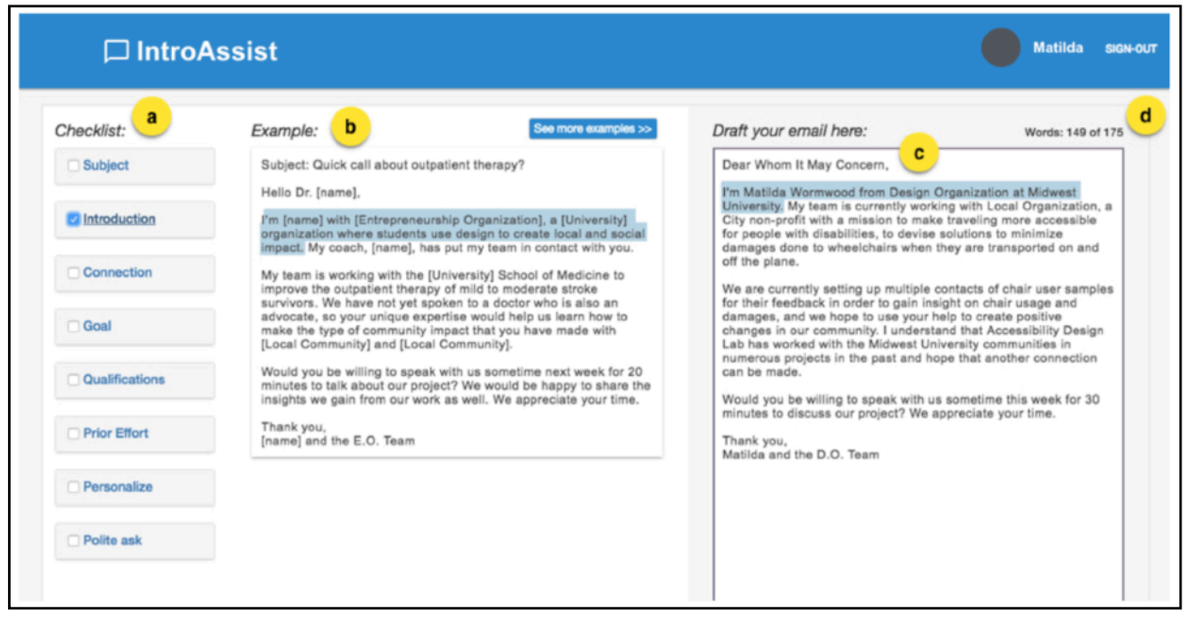

Motif

- encourages specific video structure via common patterns of video shots and a grammar to link them
- structure of form comes from excellent examples
- study focused on quality; found people could tie *capturing* and *construction*

IntroAssist

- encourages best practices with *modeling* (tagged examples), *coaching* (checklist), and *reflection* (tagging own text)
- structure of form comes from expert interviews and excellent examples
- study focused on learning; found people improved even when tool was removed

Kim, et al. “Motif: Supporting novice creativity through expert patterns.” CHI 2015.
Hui, et al. “IntroAssist: A tool to support writing introductory help requests.” CHI 2018.

generation & implementation

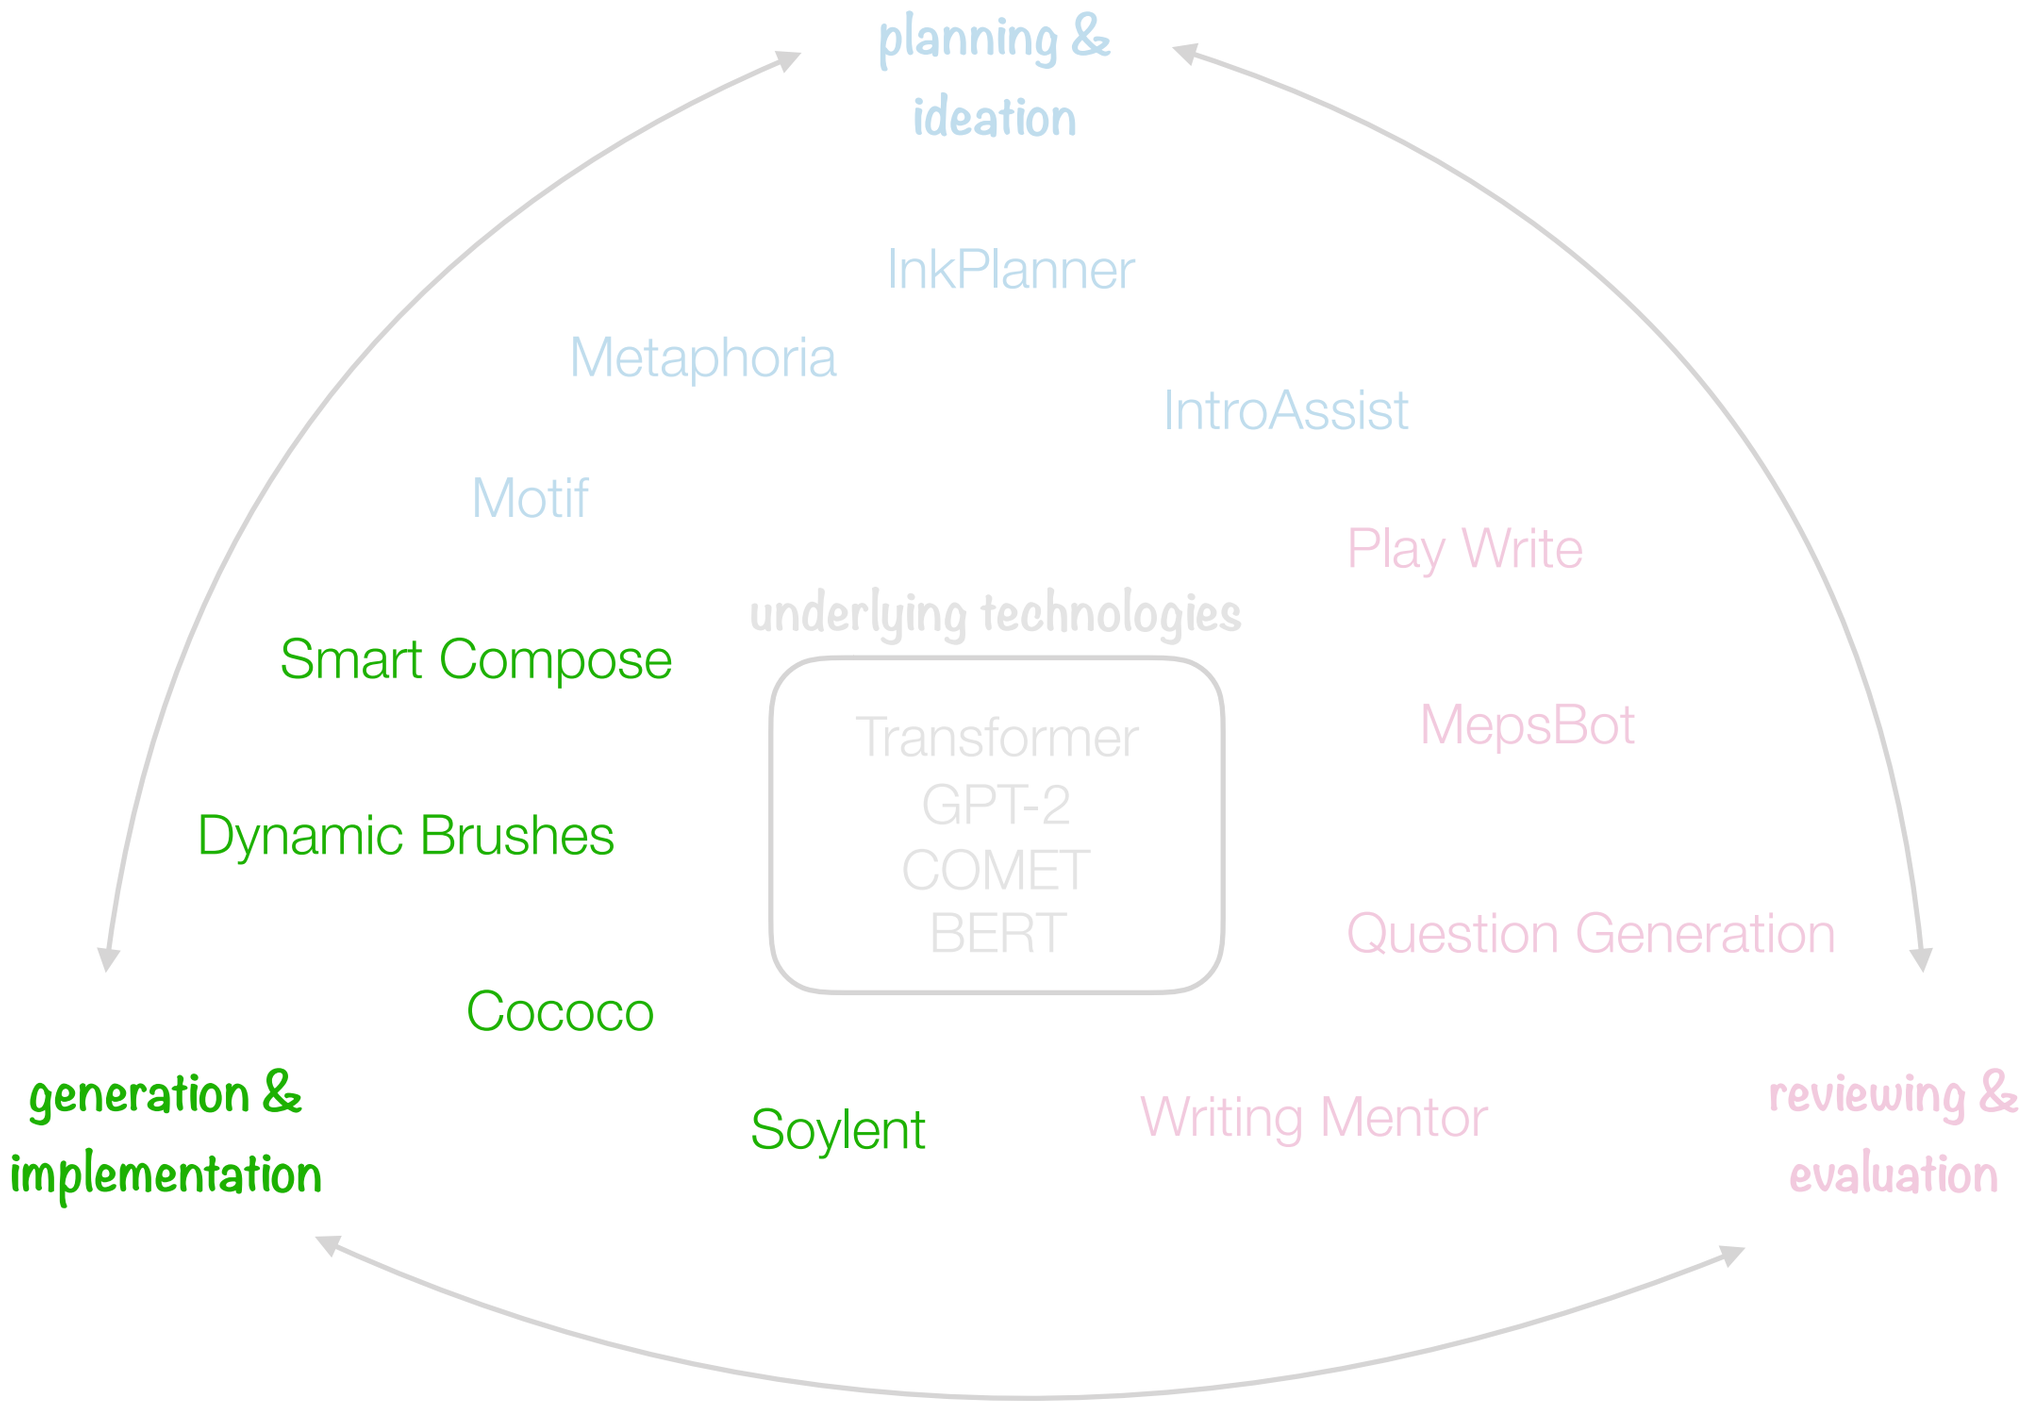

Increasing efficiency via computer or crowd.

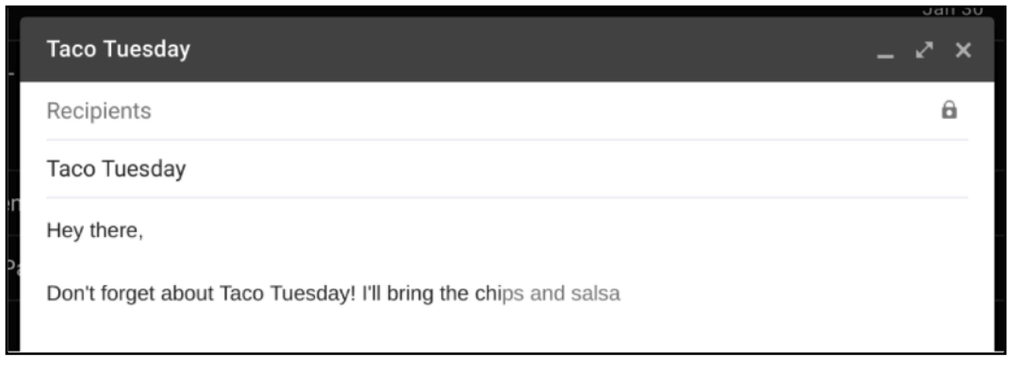

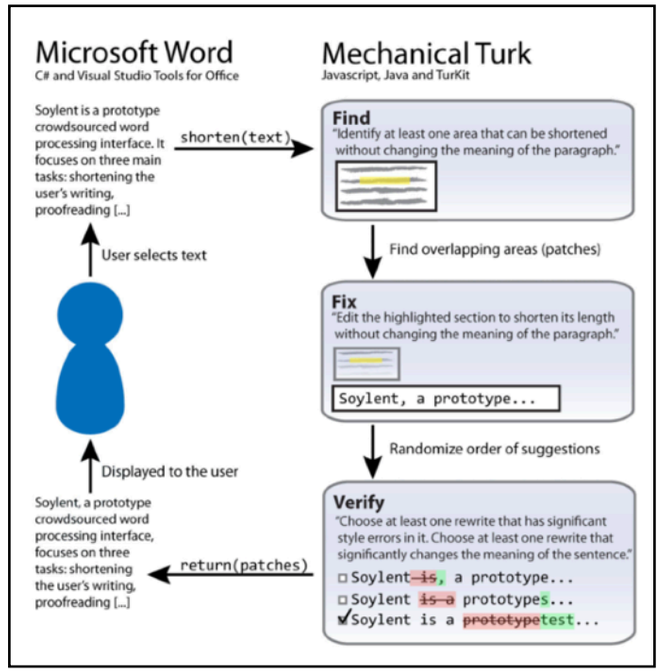

Smart Compose

- reduces repetitive typing
- suggestions come from large-scale neural language model, combined with a personalized small-scale n-gram model
- system only makes suggestions for high-probability prefixes

Soylent

- outsources editing tasks
- edits come from crowd-workers using the find-fix-verify pattern; demonstrates proof-reading and shortening tasks
- writer requests and delimits writing tasks for the workers

Chen, et al. “Gmail smart compose: Real-time assisted writing.” KDD 2019.
Bernstein, et al., “Soylent: A word processor with a crowd inside.” UIST 2010.

Again, support for process versus for semantics.

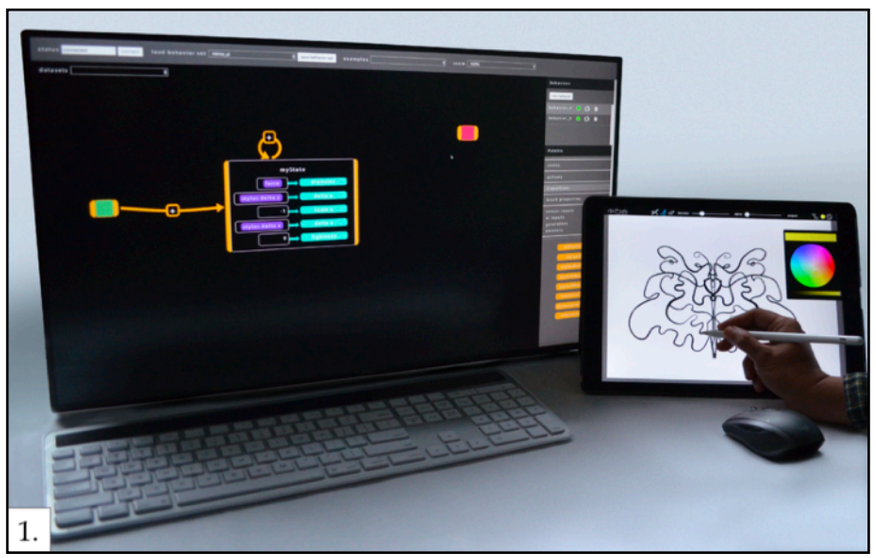

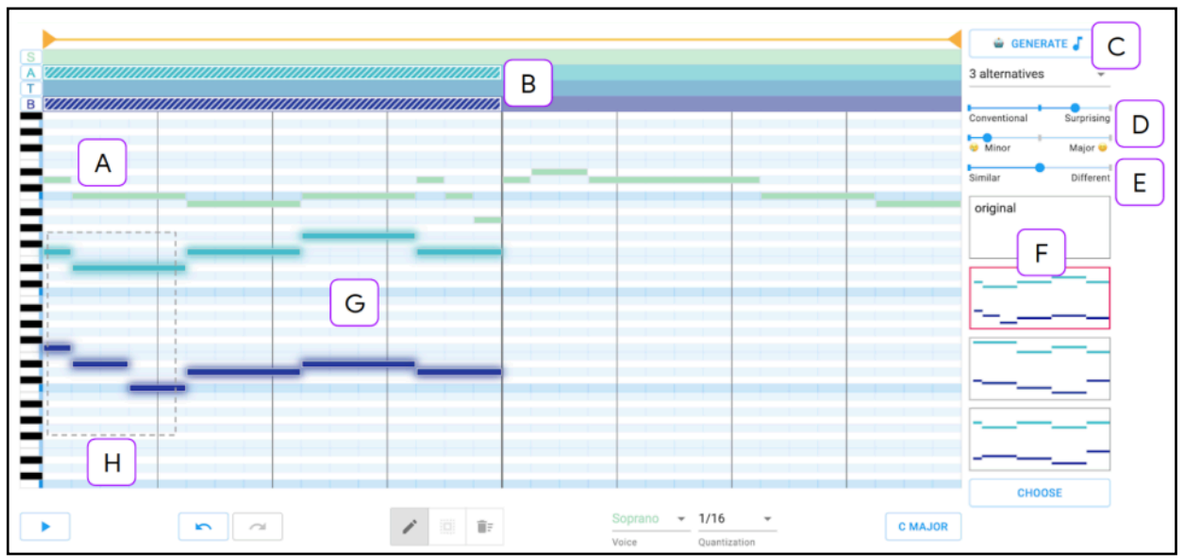

Dynamic Brushes

- integrates drawing into procedural generation with a drawing-native visual programming language
- difficult to integrate all aspects of drawing
- despite being drawing-native, still had a significant learning curve for the programming concepts

Cococo

- integrates semantic “steering" tools for generating music composition with AI; designed for novice composers
- difficult to provide all steering desires
- despite more control, still was difficult to understand, predict, or debug AI behavior
- AI introduces questions of ownership

Jacobs, et al. “Extending manual drawing practices with artist-centric programming tools.” CHI 2018. (dynamic brushes)
Louie, et al. “Novice-AI music co-creation via AI-steering tools for deep generative models.” CHI 2020. (cococo)

Support for process versus for semantics.

InkPlanner

- supports planning with better diagramming tools
- writer jots down unstructured text on tablet interface; tablet supports mind-mapping and outlining (as well as brainstorming)
- system mainly augments *process*; thus the evaluation focuses on usability

Metaphoria

- supports ideation with suggested metaphorical connections
- writer inputs seed word and requests help with button outside of text-area; tool suggests 10 potential metaphors
- system mainly augments *semantics*; thus the evaluation focuses on quality

Lu, et al. “Inkplanner: Supporting prewriting via intelligent visual diagramming.” TVCG 2018.
Gero & Chilton. “Metaphora: An algorithmic companion for metaphor creation.” CHI 2019.

Again, support for process versus for semantics.

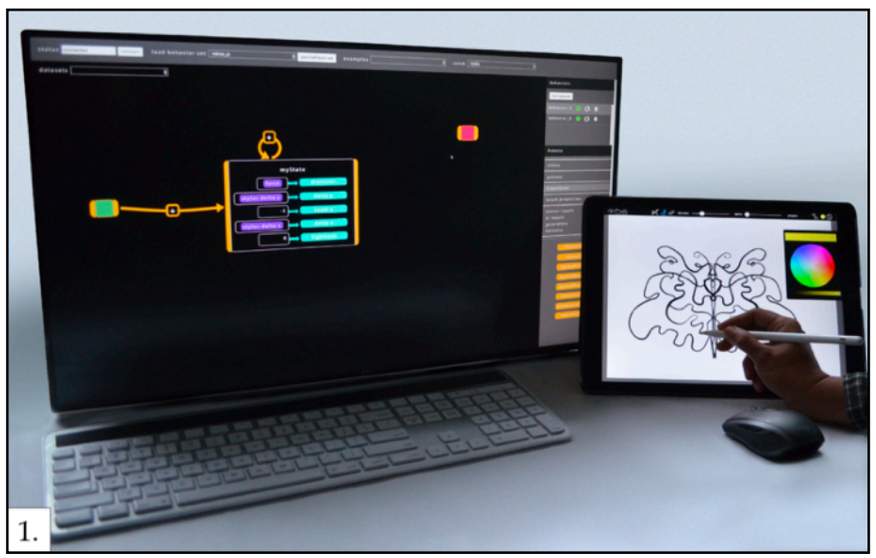

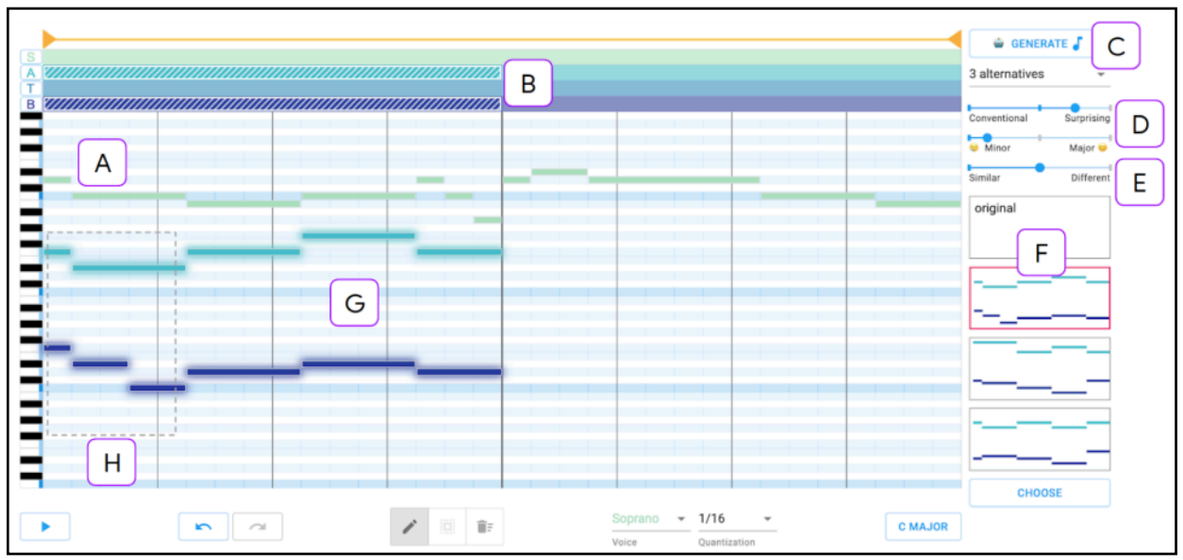

Dynamic Brushes

- integrates drawing into procedural generation with a drawing-native visual programming language
- difficult to integrate all aspects of drawing
- despite being drawing-native, still had a significant learning curve for the programming concepts

Cococo

- integrates semantic “steering" tools for generating music composition with AI; designed for novice composers
- difficult to provide all steering desires
- despite more control, still was difficult to understand, predict, or debug AI behavior
- AI introduces questions of ownership

Jacobs, et al. “Extending manual drawing practices with artist-centric programming tools.” CHI 2018. (dynamic brushes)
Louie, et al. “Novice-AI music co-creation via AI-steering tools for deep generative models.” CHI 2020. (cococo)

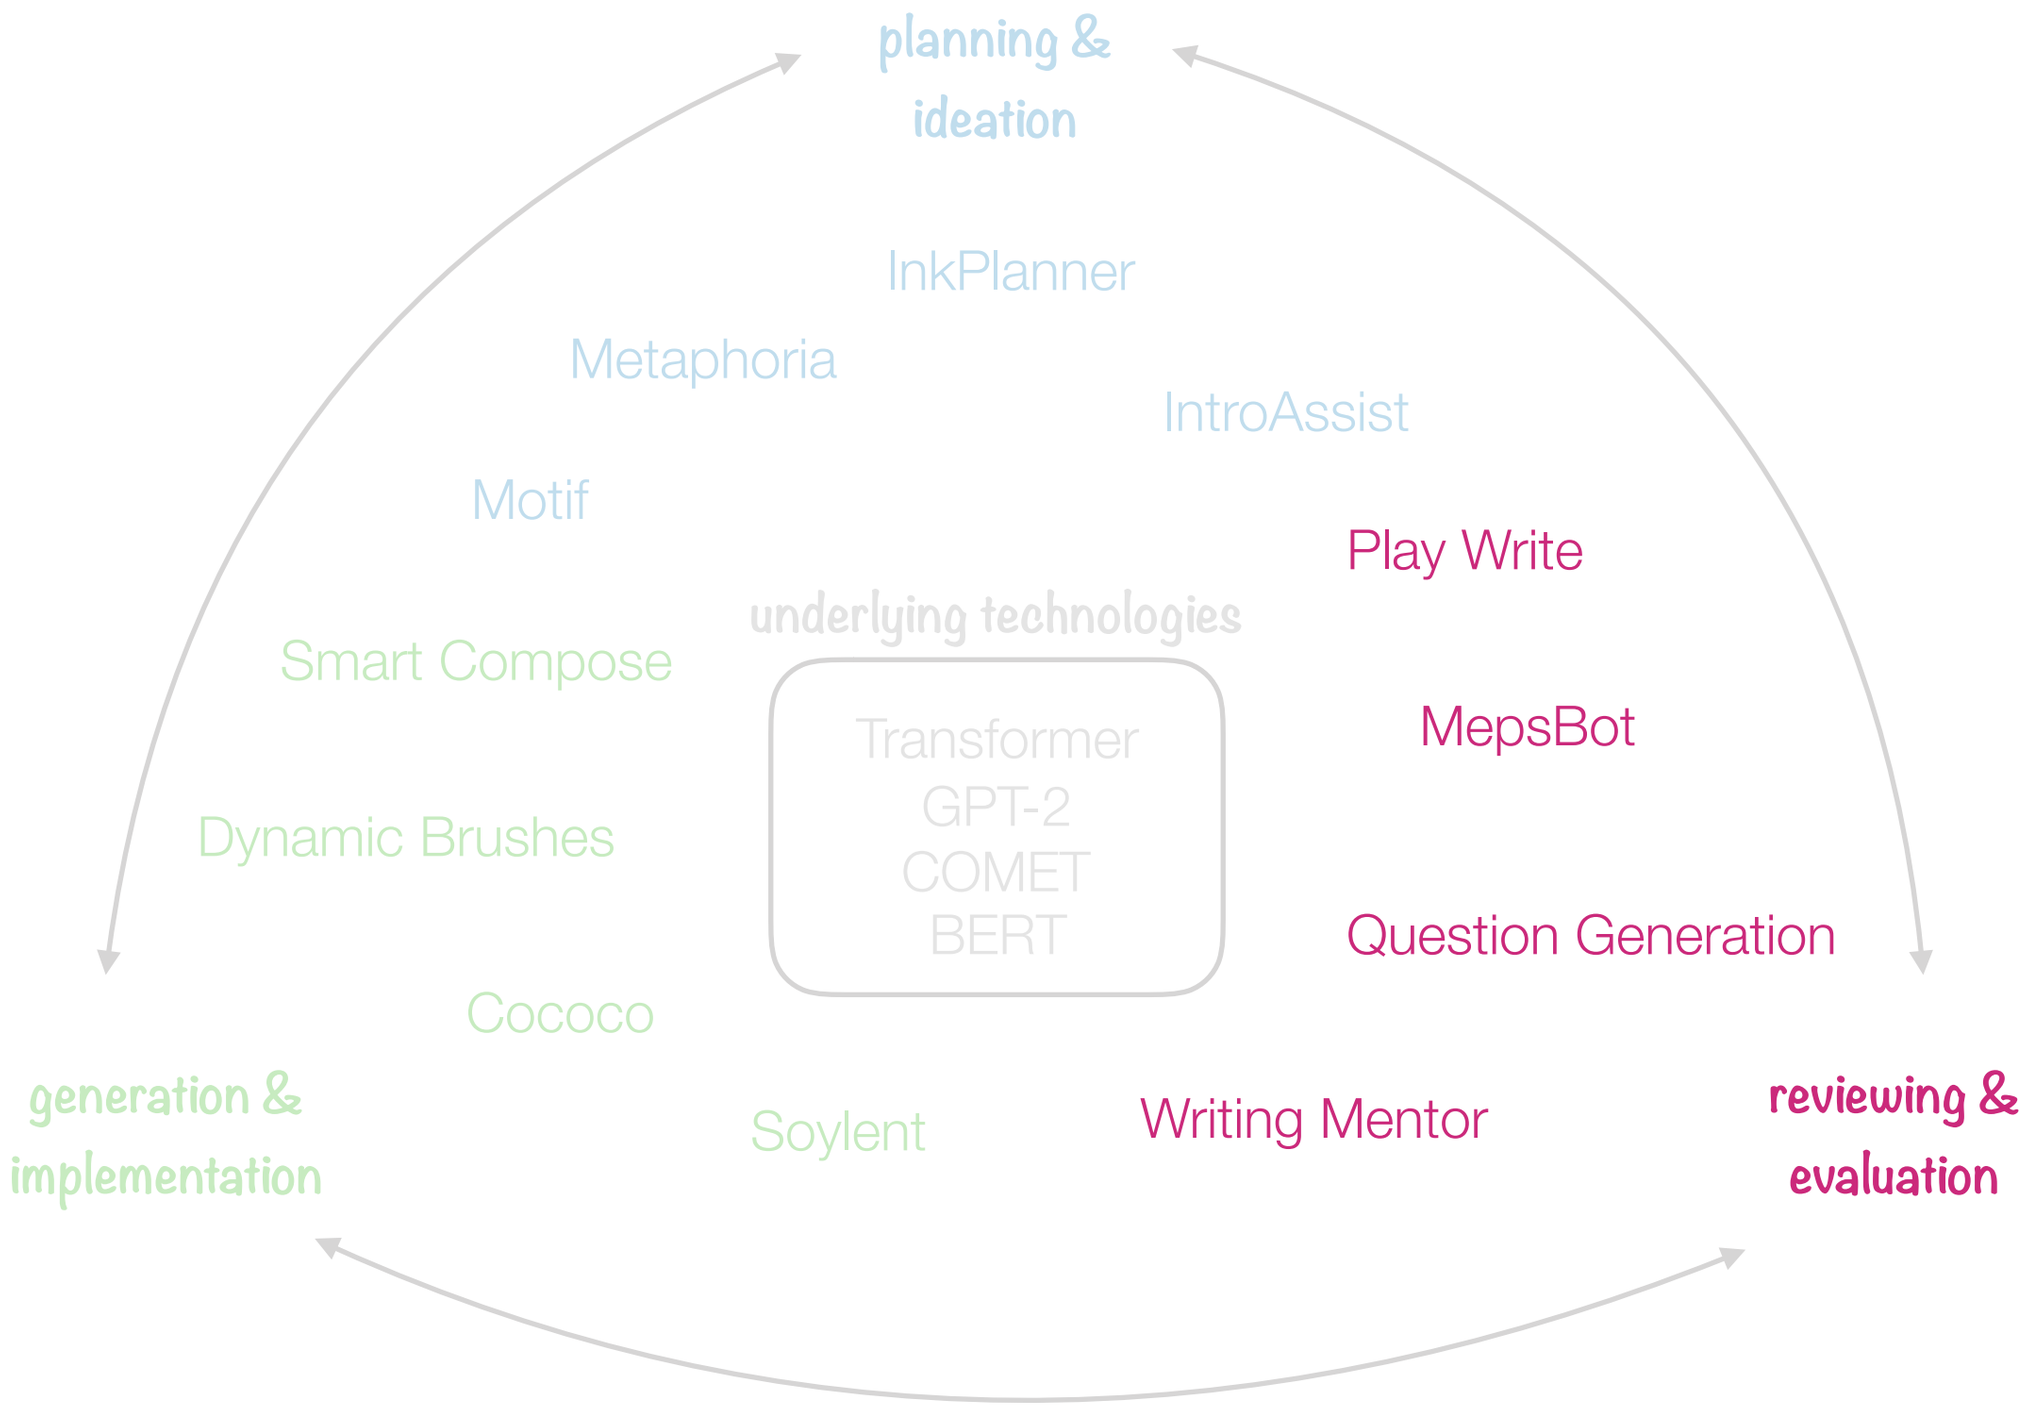
planning & ideation
InkPlanner
Metaphoria
IntroAssist
Motif
Play Write
underlying technologies
Smart Compose
Transformer
GPT-2
COMET
BERT
MepsBot
Dynamic Brushes
Question Generation
Cococo
generation & implementation
Soylent
Writing Mentor
reviewing & evaluation

Support for efficiency versus for improved quality.

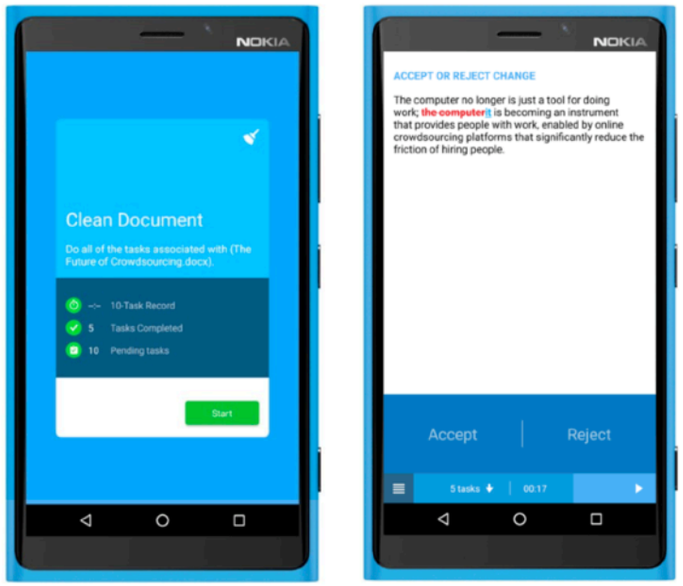

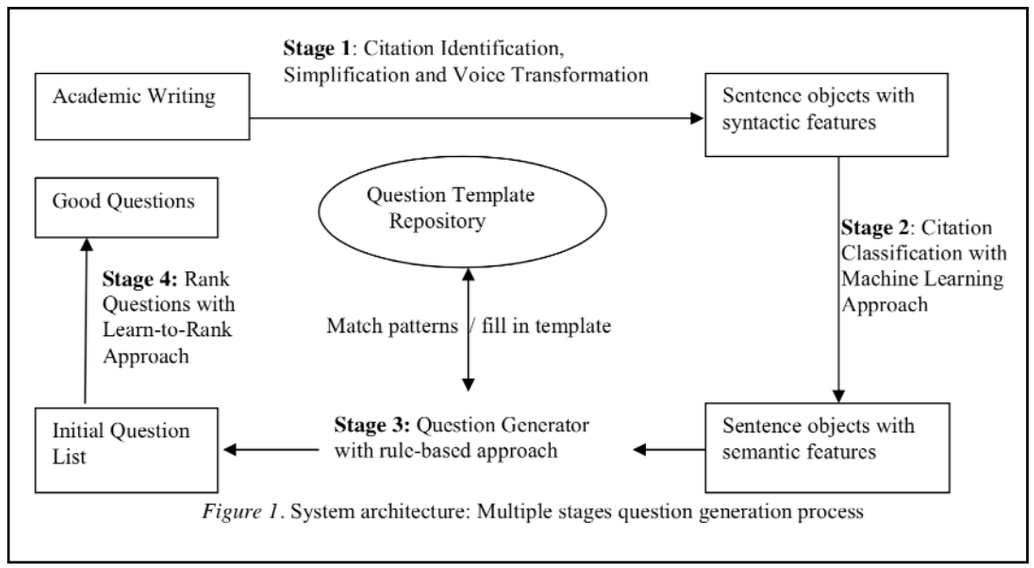

Figure 1. System architecture: Multiple stages question generation process

Play Write

- system creates review micro-tasks, e.g. correcting spelling, and identifying wordy sentences
- goal is to improve use of writer’s time
- seeks to create and transfer review tasks that can be done while multi-tasking

Question Generation

- system creates review questions, e.g. what is the research question, and is the analysis of data accurate
- goal is to encourage further reflection
- seeks to replicate the review process of teachers and other students

Iqbal, et al. “Multitasking with Play Write, a model microproductivity writing tool.” UIST 2018.
Liu, et al., “Automatic generation and ranking of questions for critical review.” Journal of Edu. Technology & Society 2014

Replicating human review with automatic assessment.

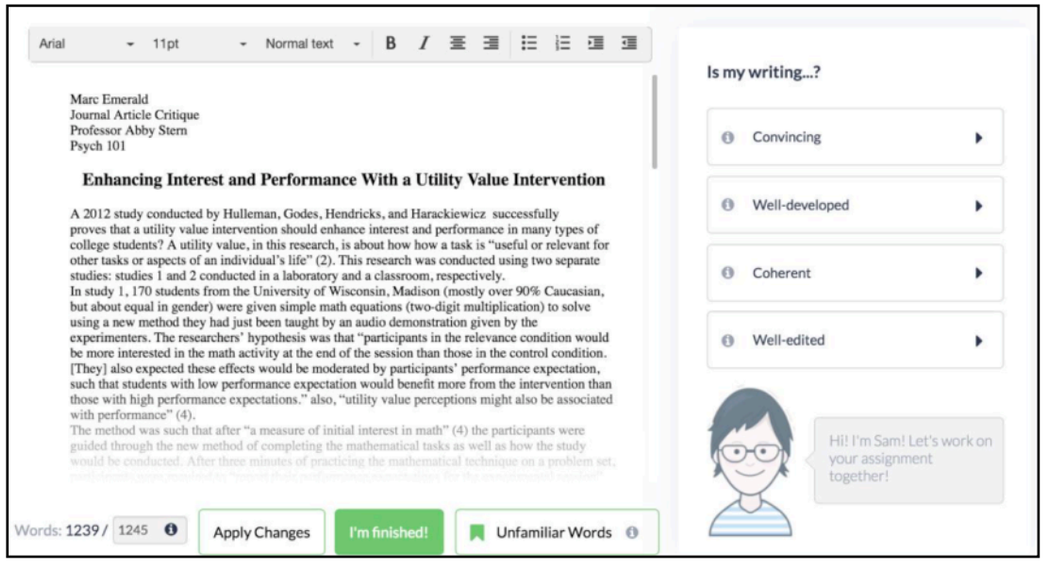

Mental Peer Support Bot (MepsBot)

- provides assessment and recommendations on informational and emotional support, as well as related examples
- support is based on the specific qualities of comments in mental health forums
- builds on existing NLP technologies using custom data set

Writing Mentor

- provides actionable feedback on being *convincing*, *well-developed*, *coherent*, and *well-edited*
- support is for general postsecondary writing, and thus contains many features
- builds on a variety of existing NLP technologies

Peng, et al. “Exploring the effects of technological writing assistance for support providers in online…” CHI 2020. (mepsbot)
Burstein, et al. “Writing mentor: Writing progress using self-regulated writing support.” Journal of Writing Analytics 2019.

2. The “process” view of writing support.

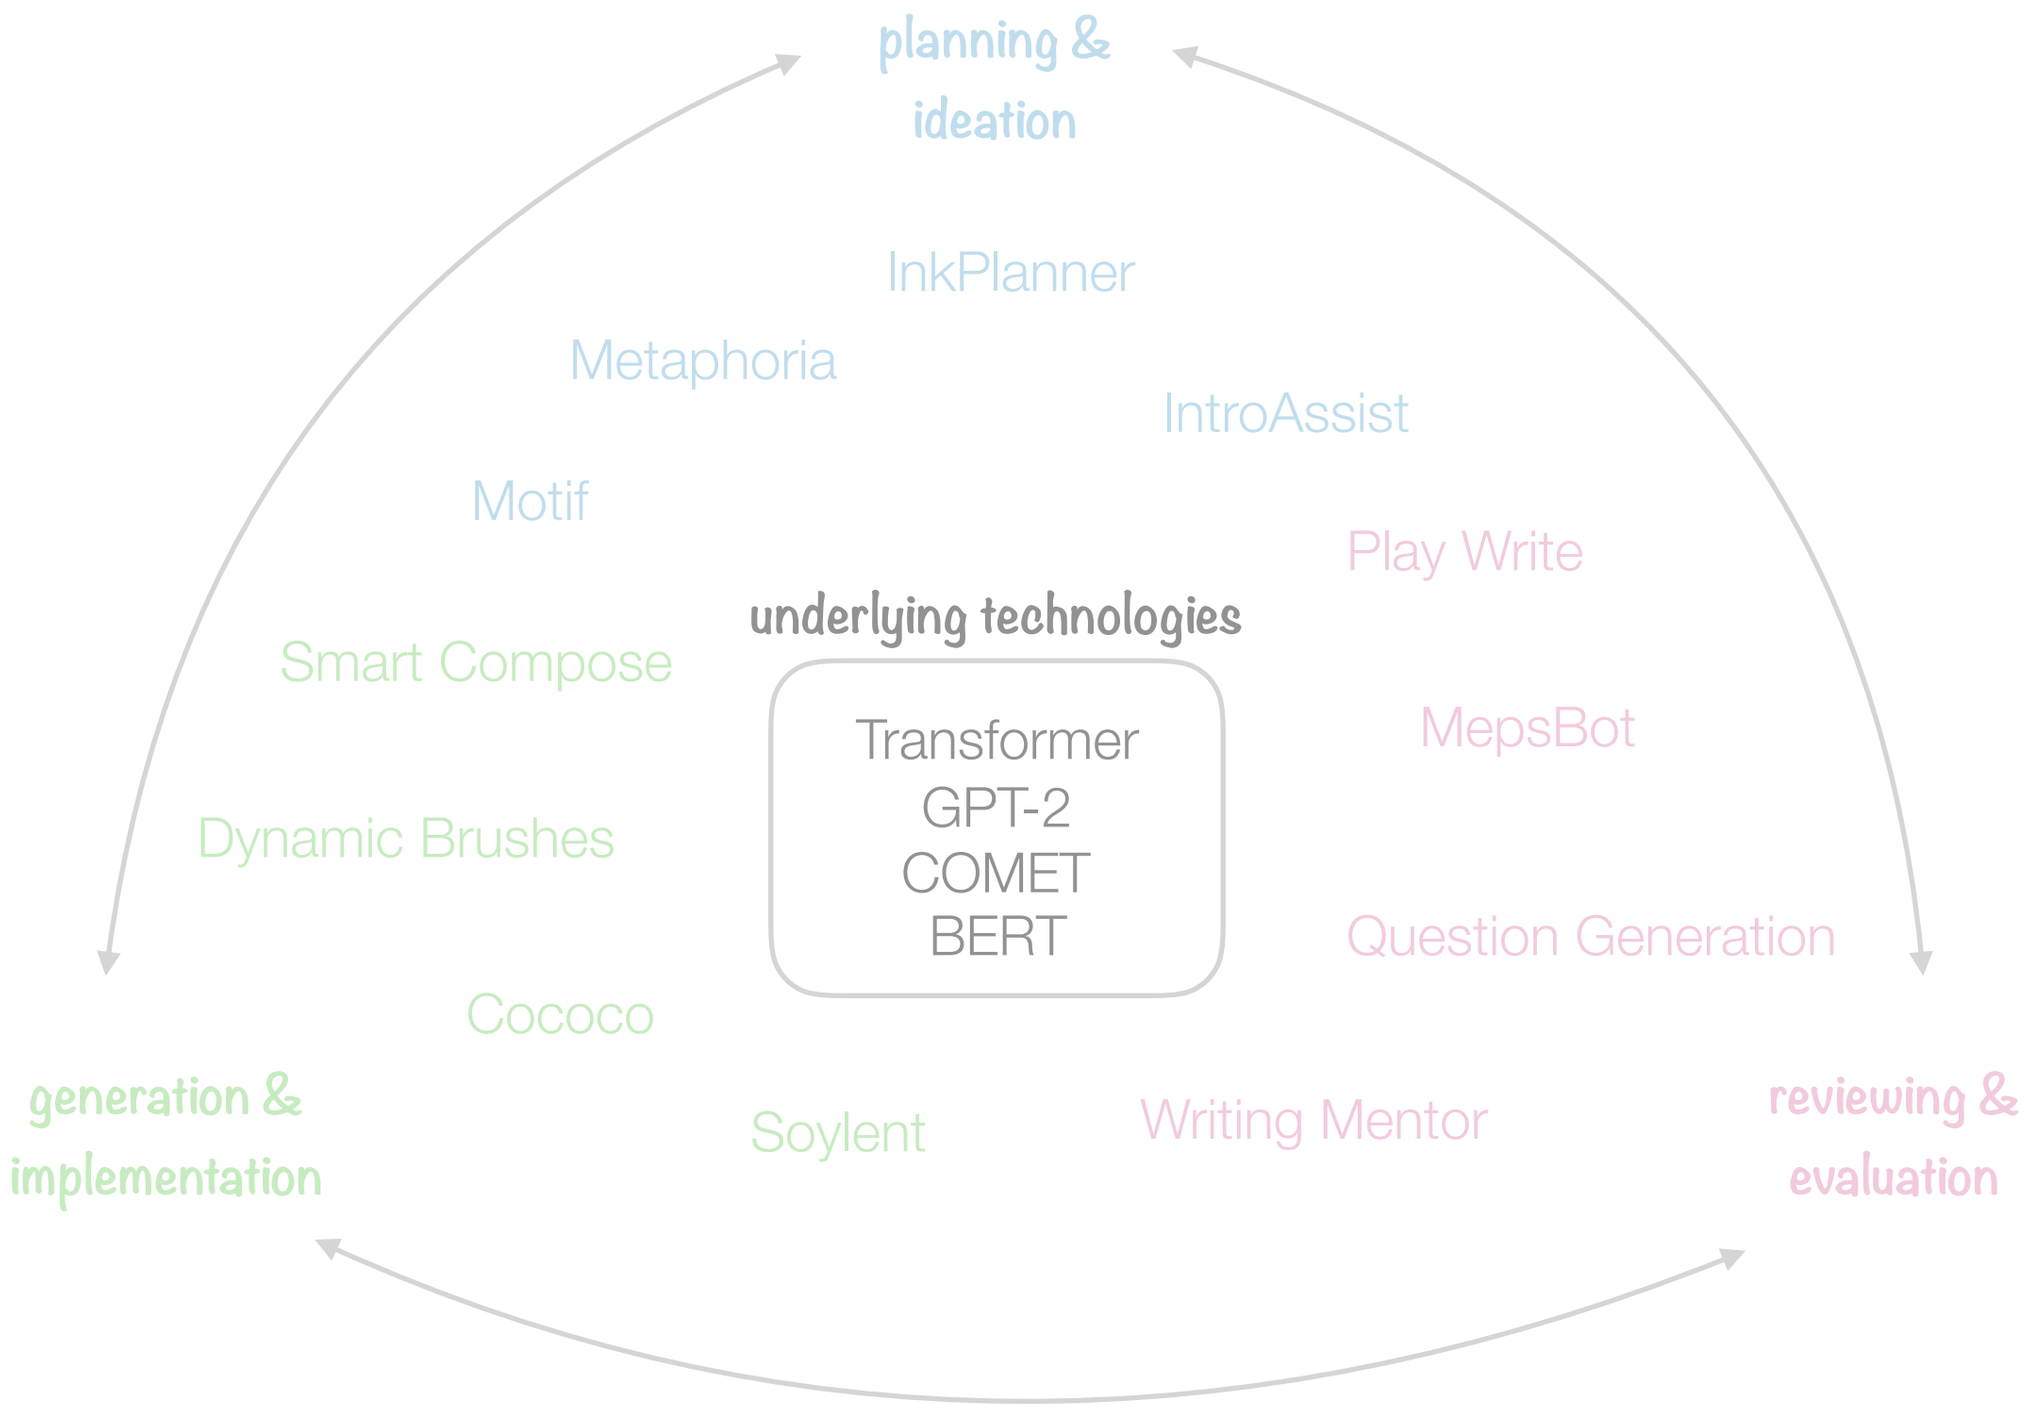

Where are underlying technologies currently used?

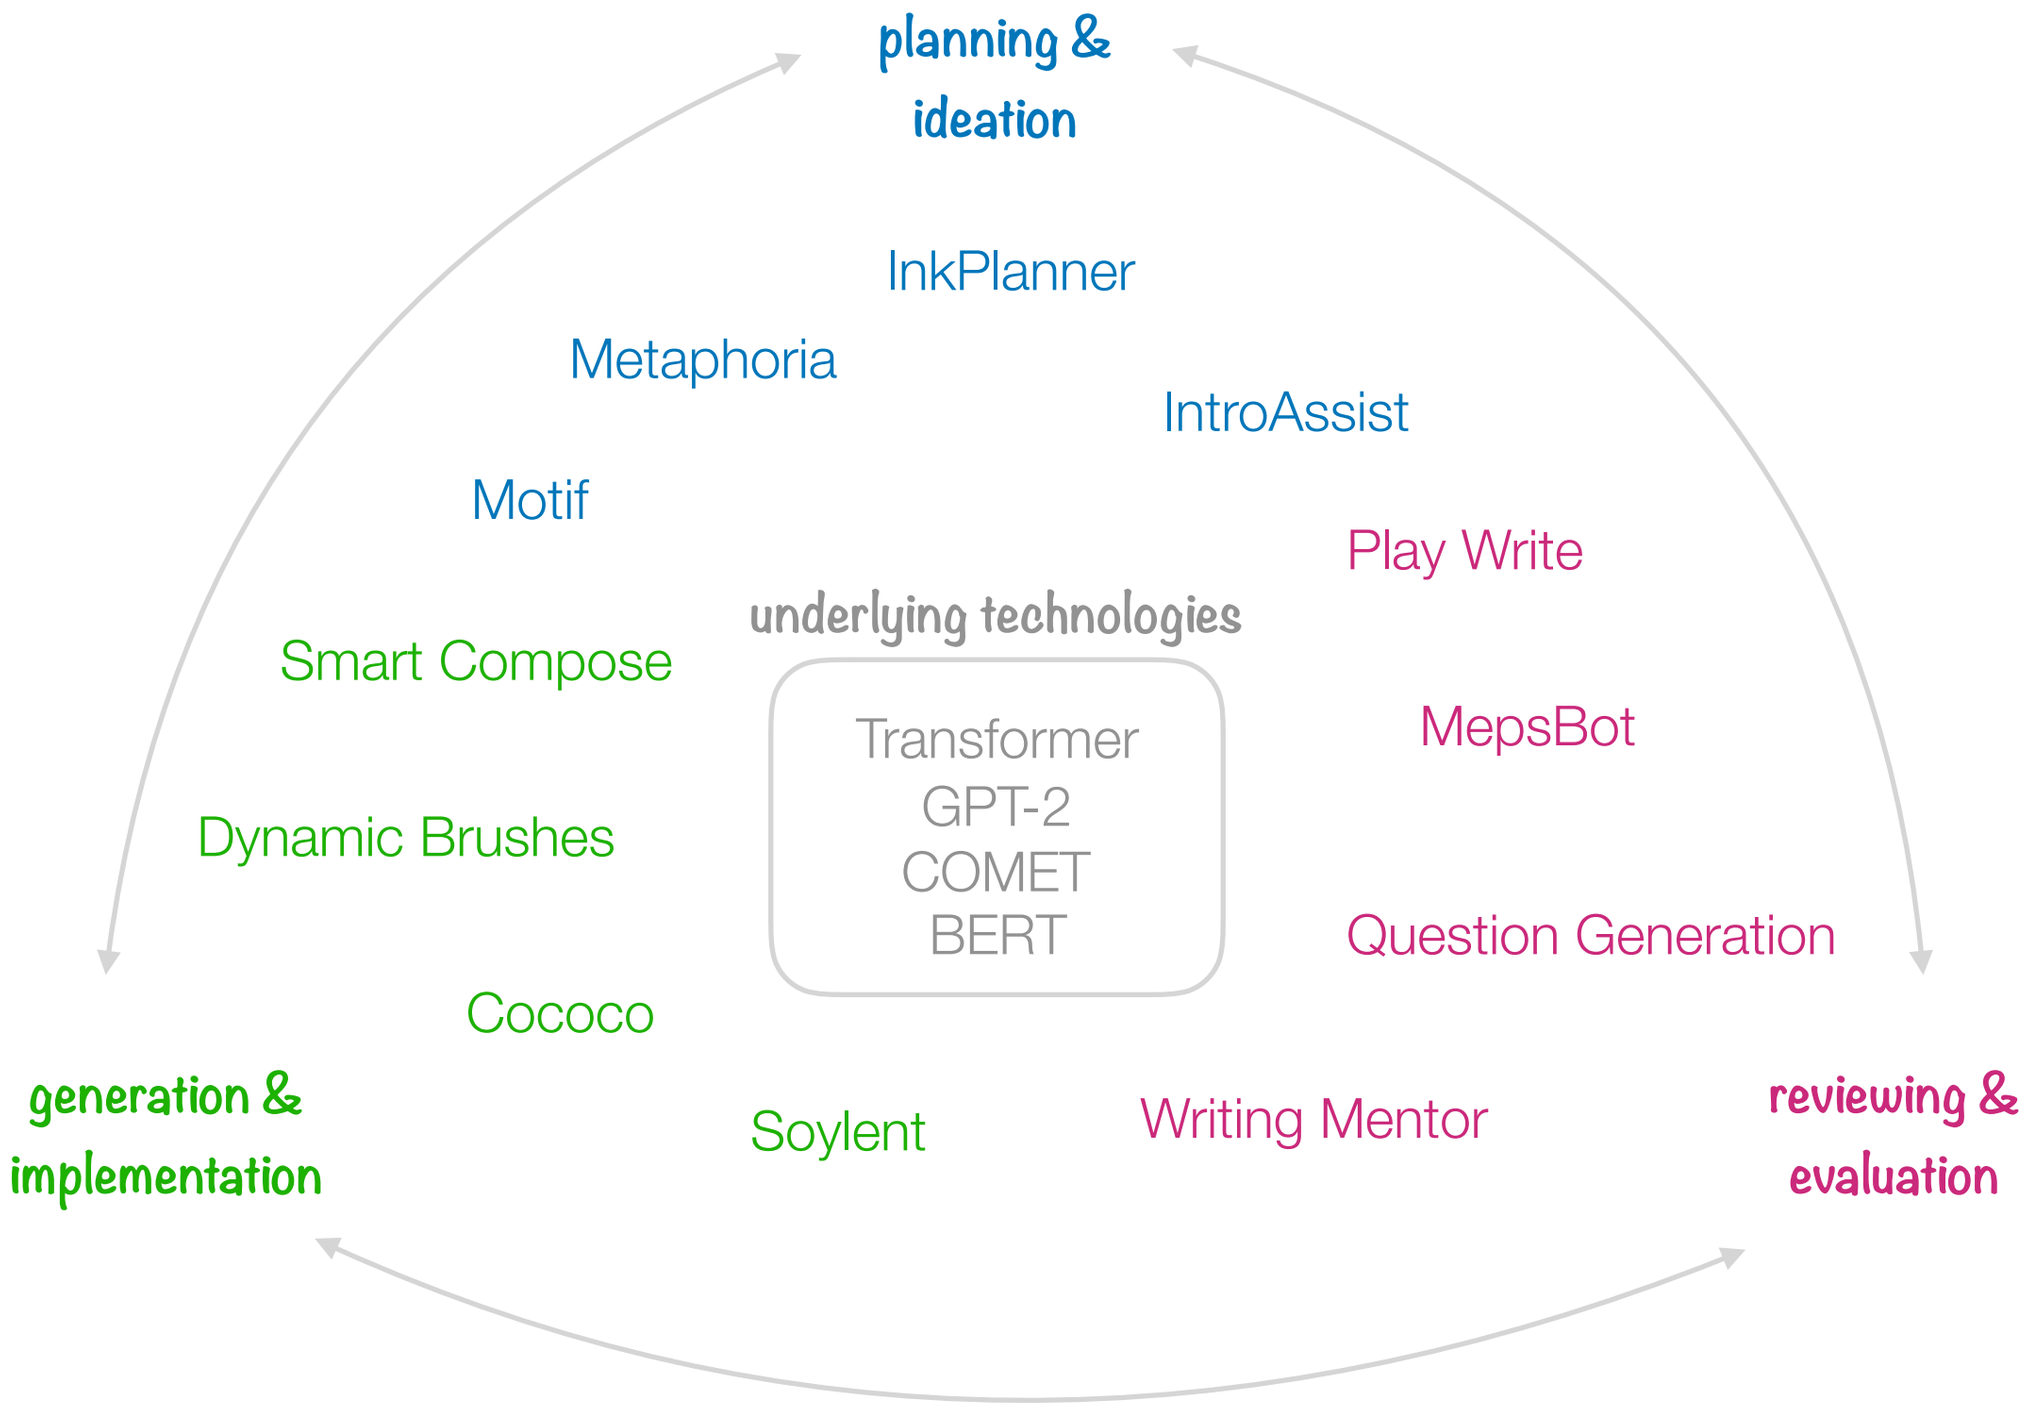

Where are underlying technologies currently used?

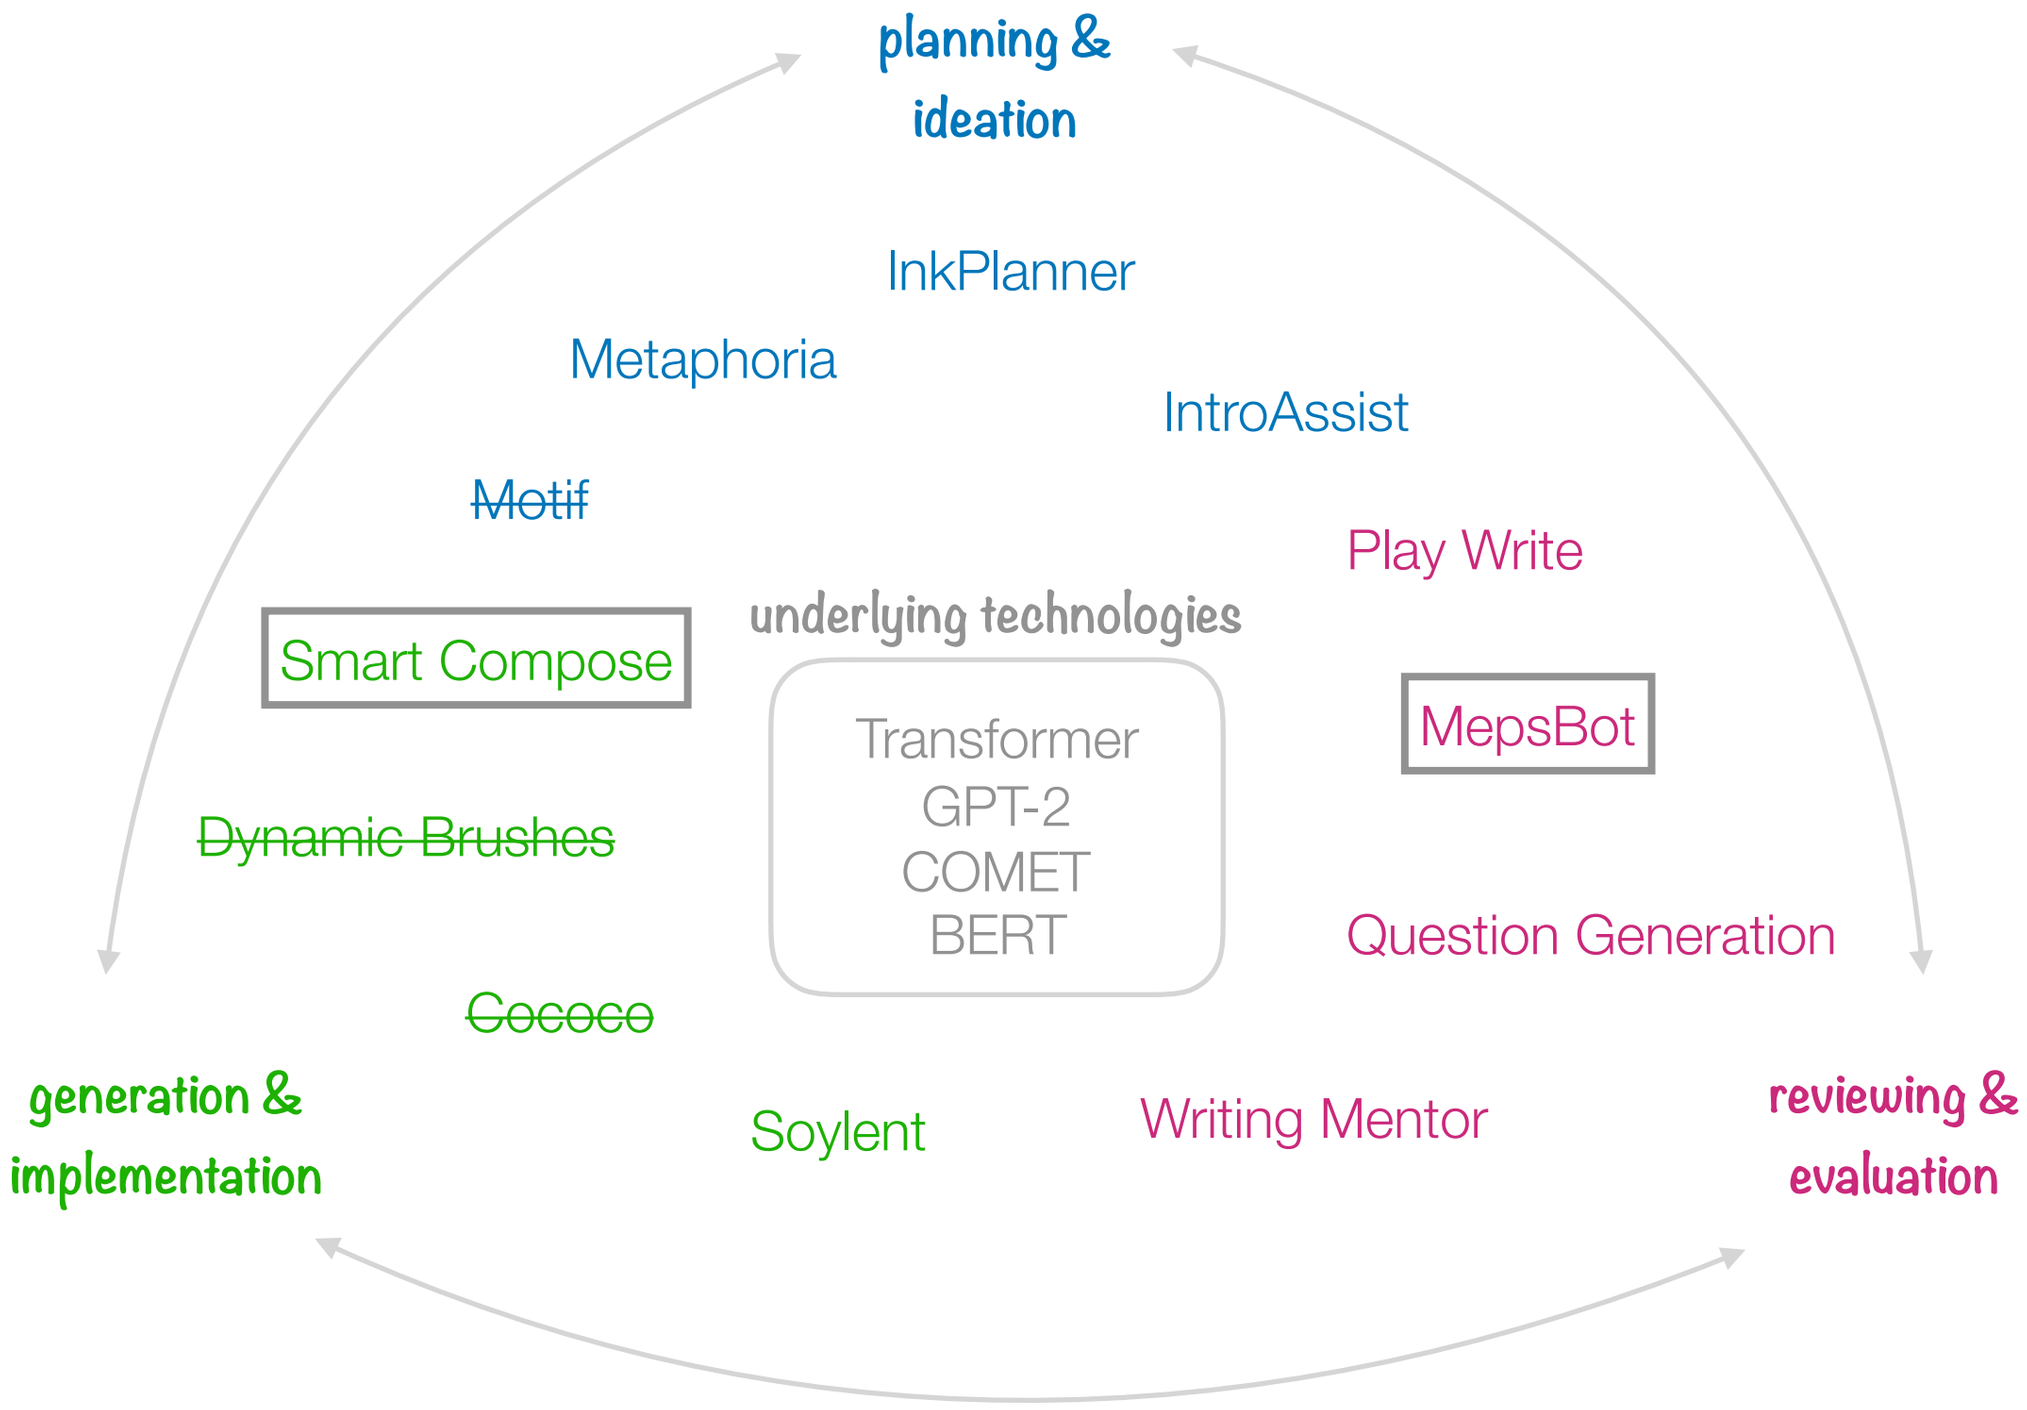

Language models assign probabilities to sequences.

Language Models

- model assigns probabilities to sequences of words
- n-gram is the simplest
- uses:
  - speech recognition
  - grammar checking
  - machine translation
  - conditional generation

Recurrent Neural Networks

- neural language model that uses a hidden state to hold contextual information
- several variants to improve managing long-term context: LSTMs and GRUs

Transformer Model

- multi-headed self-attention
- improves upon RNNs in:
  - modeling longer-term dependencies
  - decreasing training time (no recurrence; self-attention is a big matrix)
- interpreting model outputs

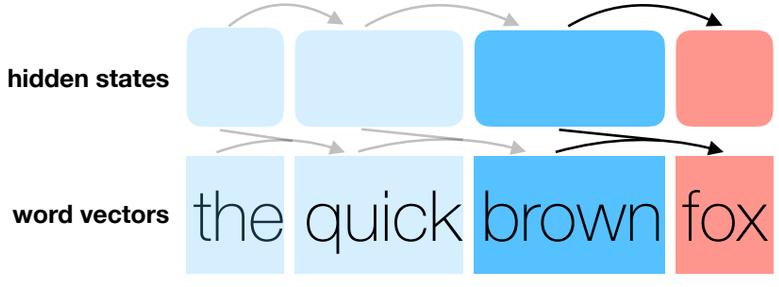

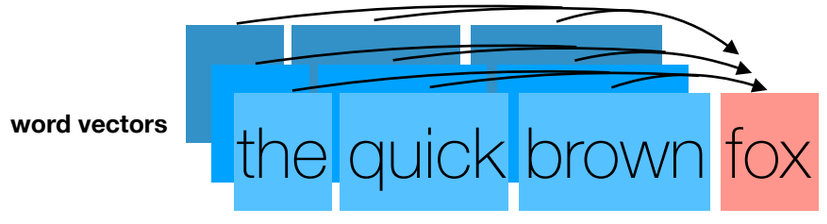

Vaswani, et al. “Attention is all you need.” NeurIPS 2017. (transformer model)

Language models as multi-purpose tools.

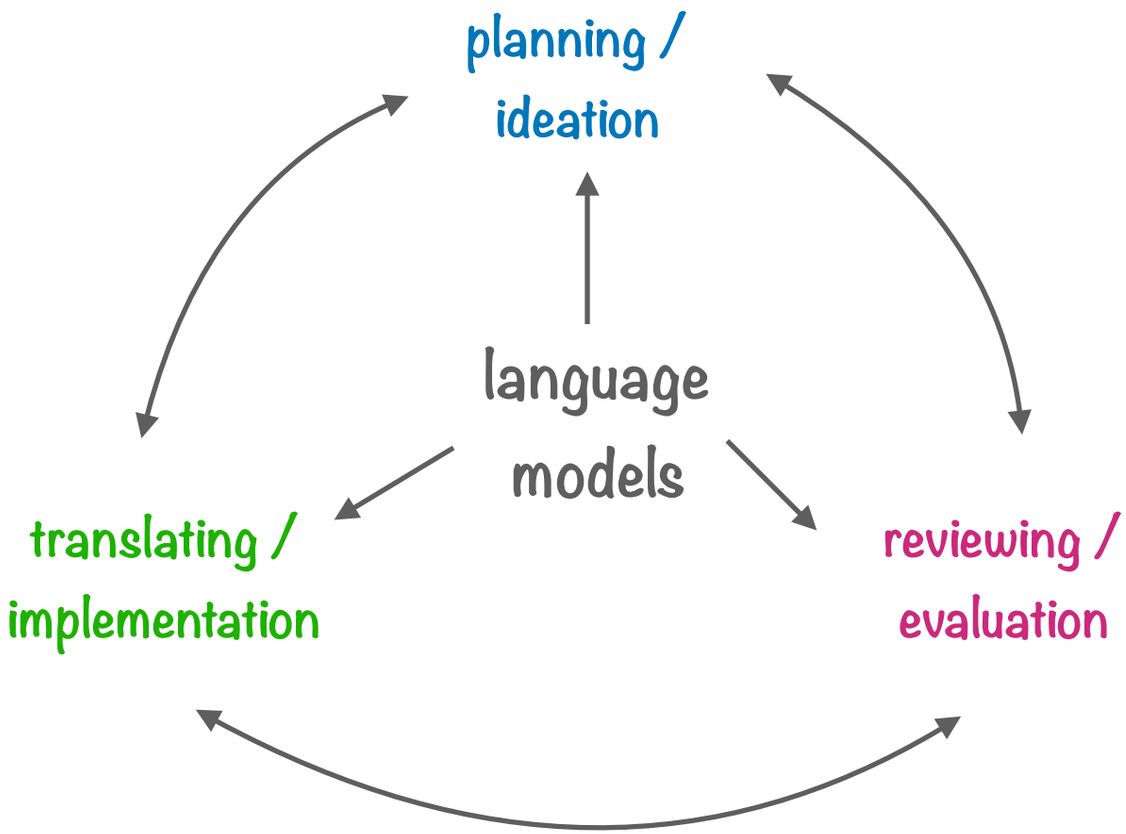

Language models as generative tools.

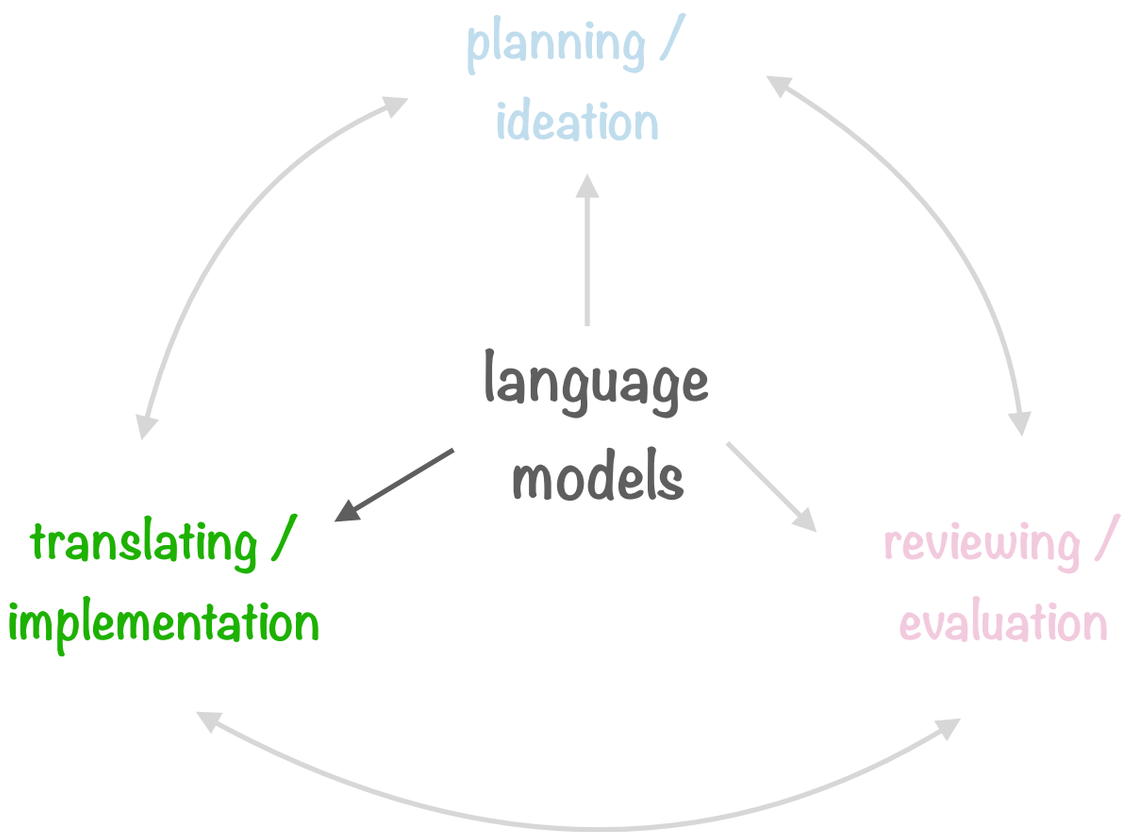

GPT-2

- uses transformer model with a context of 1024 tokens (byte-pair encoding)
- training on WebText corpus (40GB)
- no fine-tuning for downstream tasks, instead simply conditions on text input
- GPT-2 does extremely well on small datasets created to measure long-term dependencies (e.g. Children’s Book and LAMBADA)
- it does less well on more complex tasks, like summarization or question answering
- it is able to generate coherent long-form texts

Radford, et al. “Language models are unsupervised multitask learners.” OpenAI 2019.

Language models drive language evaluation.

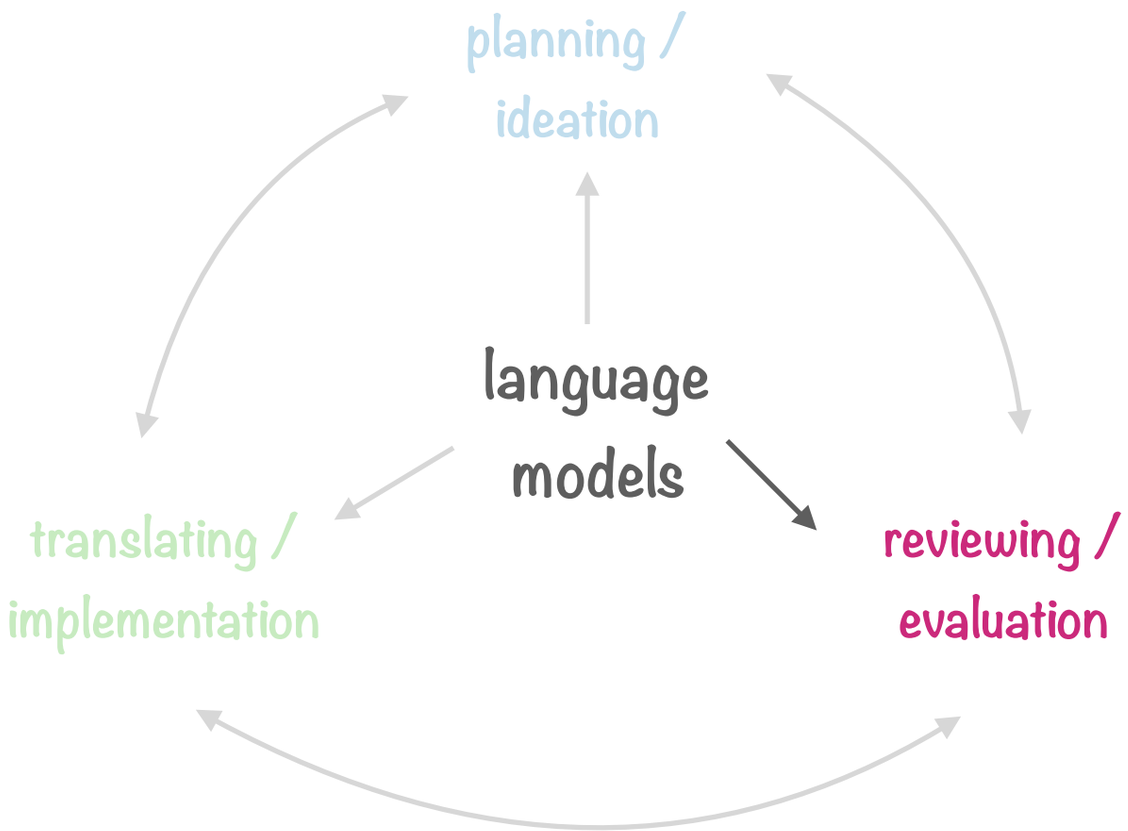

BERT

- uses the transformer language model with bi-directional pre-training: recover masked tokens and predict if 2nd sentence is real.
- unique input structure: [CLS] sent1 [SEP] sent2. <= 512 tokens.
- pre-training done with BooksCorpus (800M words) and English Wikipedia (2.5B words) and fine-tuned on downstream task data
- outperforms SOTA on General Language Understanding Evaluation (GLUE) which includes tasks like entailment and sentiment classification, and semantic similarity.
- outperforms SOTA on Stanford Question Answering Dataset (SQuAD) and more

Devlin, et al. “Bert: Pre-training of deep bidirectional transformers for language understanding.” arXiv preprint 2018.

Language models may be able to support ideation.

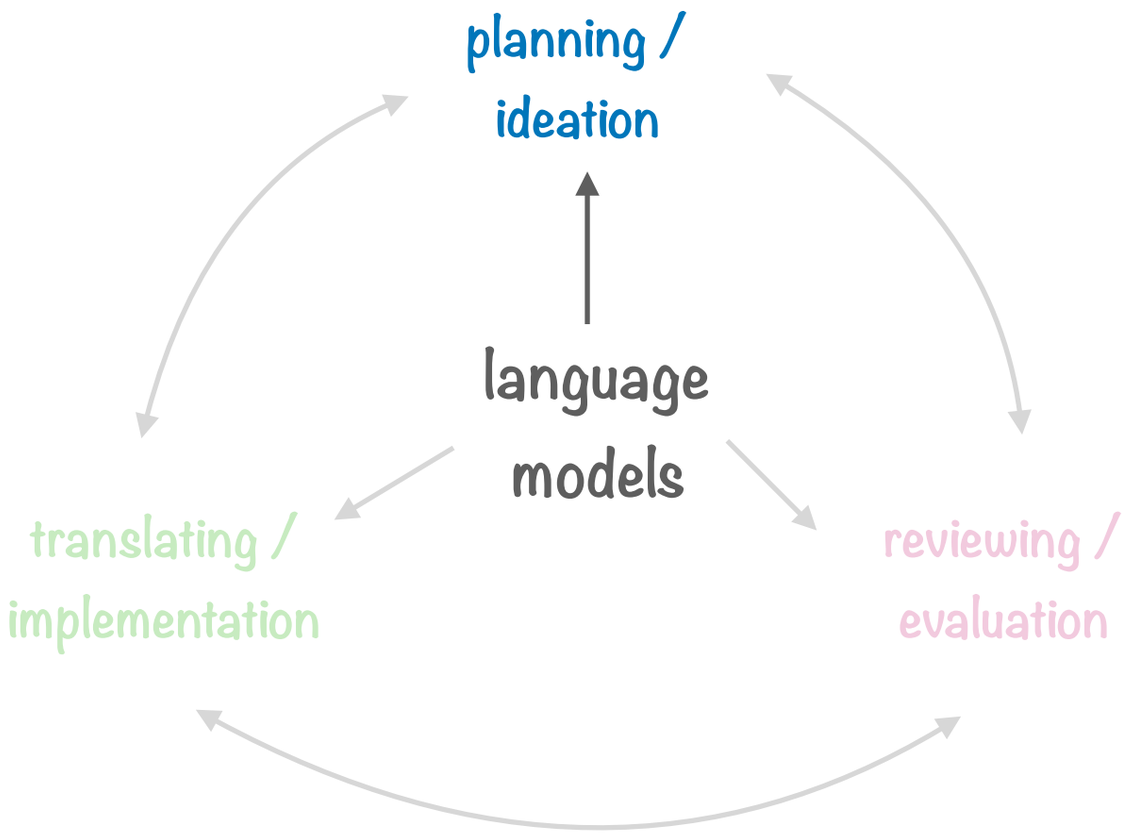

COMET

- uses pre-trained GPT transformer model (trained on BooksCorpus)
- fine-tunes on knowledge-base relations (ATOMIC and ConceptNet; input is subject and relation, output is object)
- tested by its ability to complete to knowledge-base relations in a test set that were not included in fine-tuning
- was additionally able to generate relations that weren’t in the knowledge-bases at all, demonstrating it was extracting new knowledge from the language model

Bosselut, et al. “COMET: Commonsense transformers for knowledge graph construction.” ACL 2019.

What is works and what is lacking for each process?

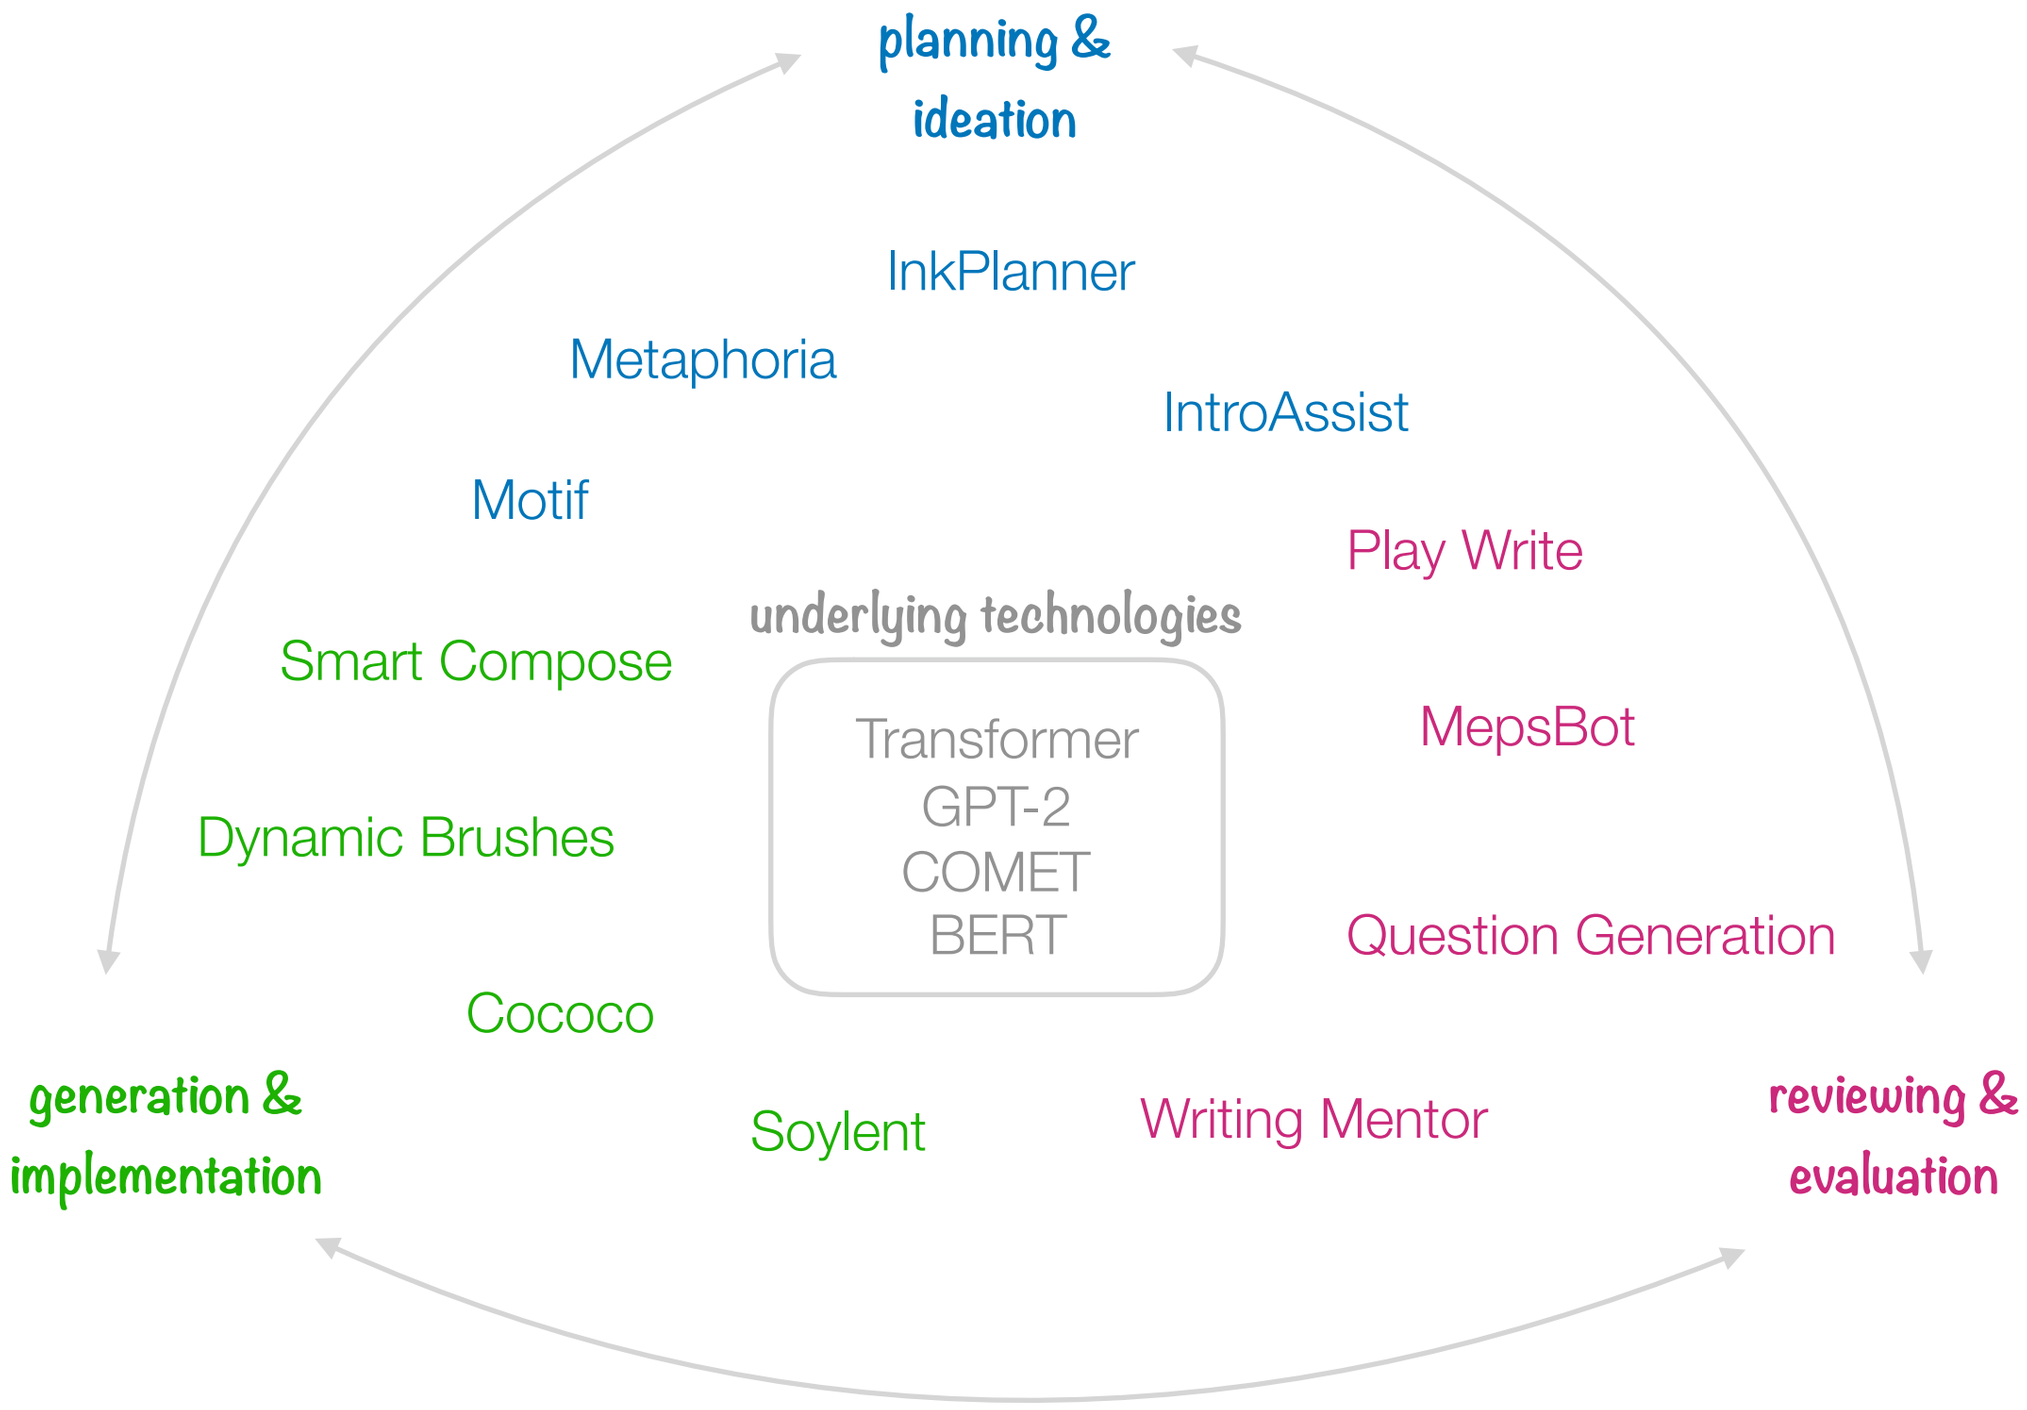

Right now, planning benefits most from expert structures. Ideation tools often fail to be semantically coherent.

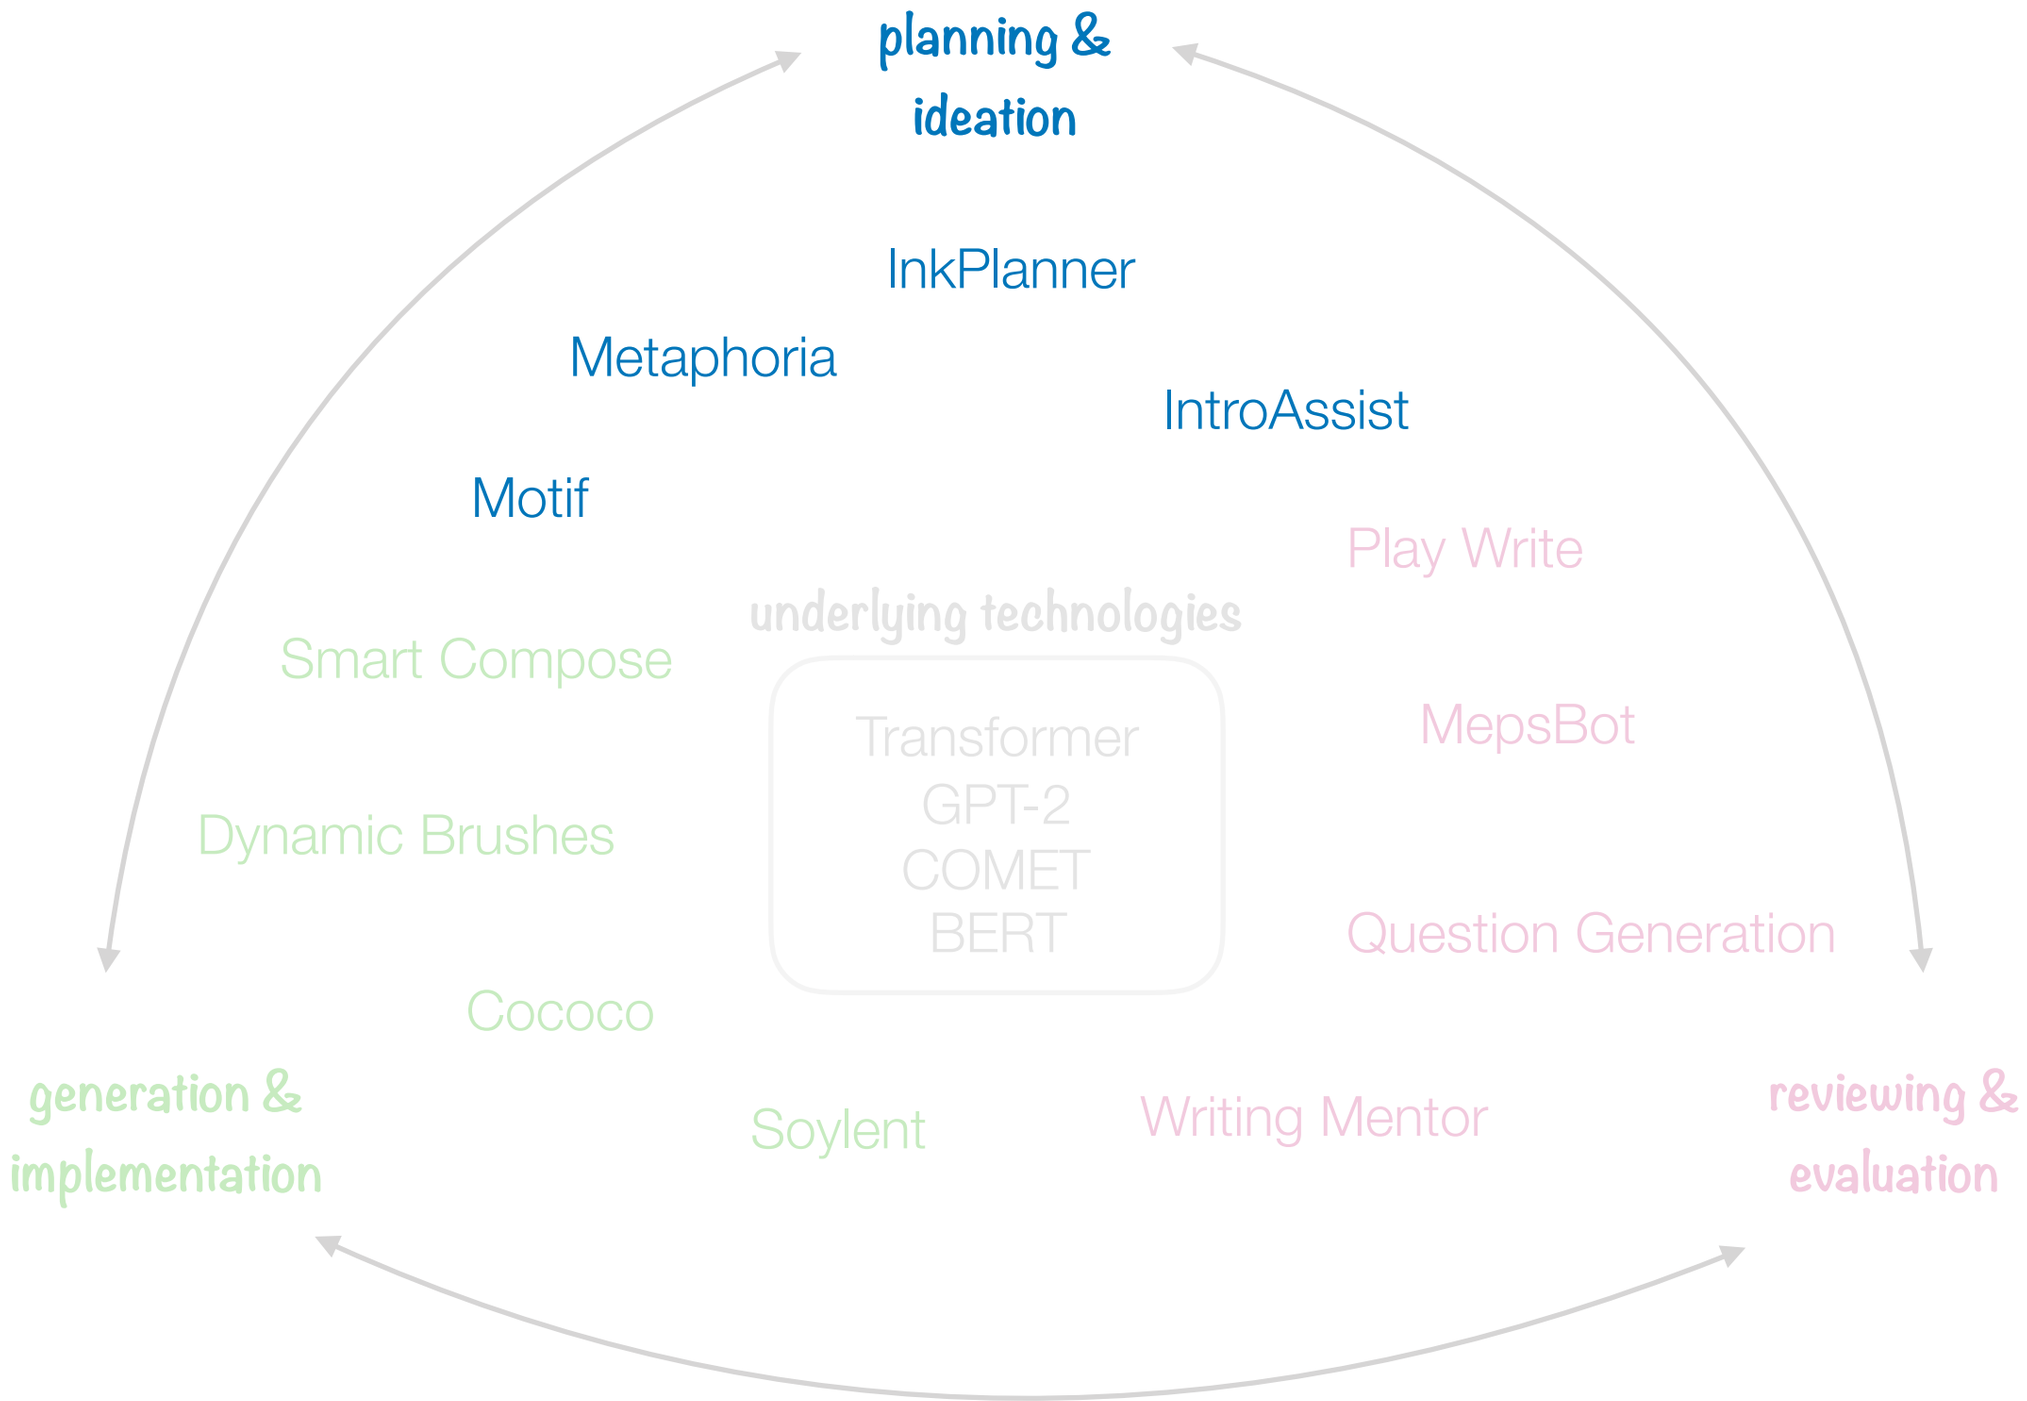

2. The “process” view of writing support.

Generative tools have promise but require new interfaces that allow people to explore the conceptual space with control, like steering tools.

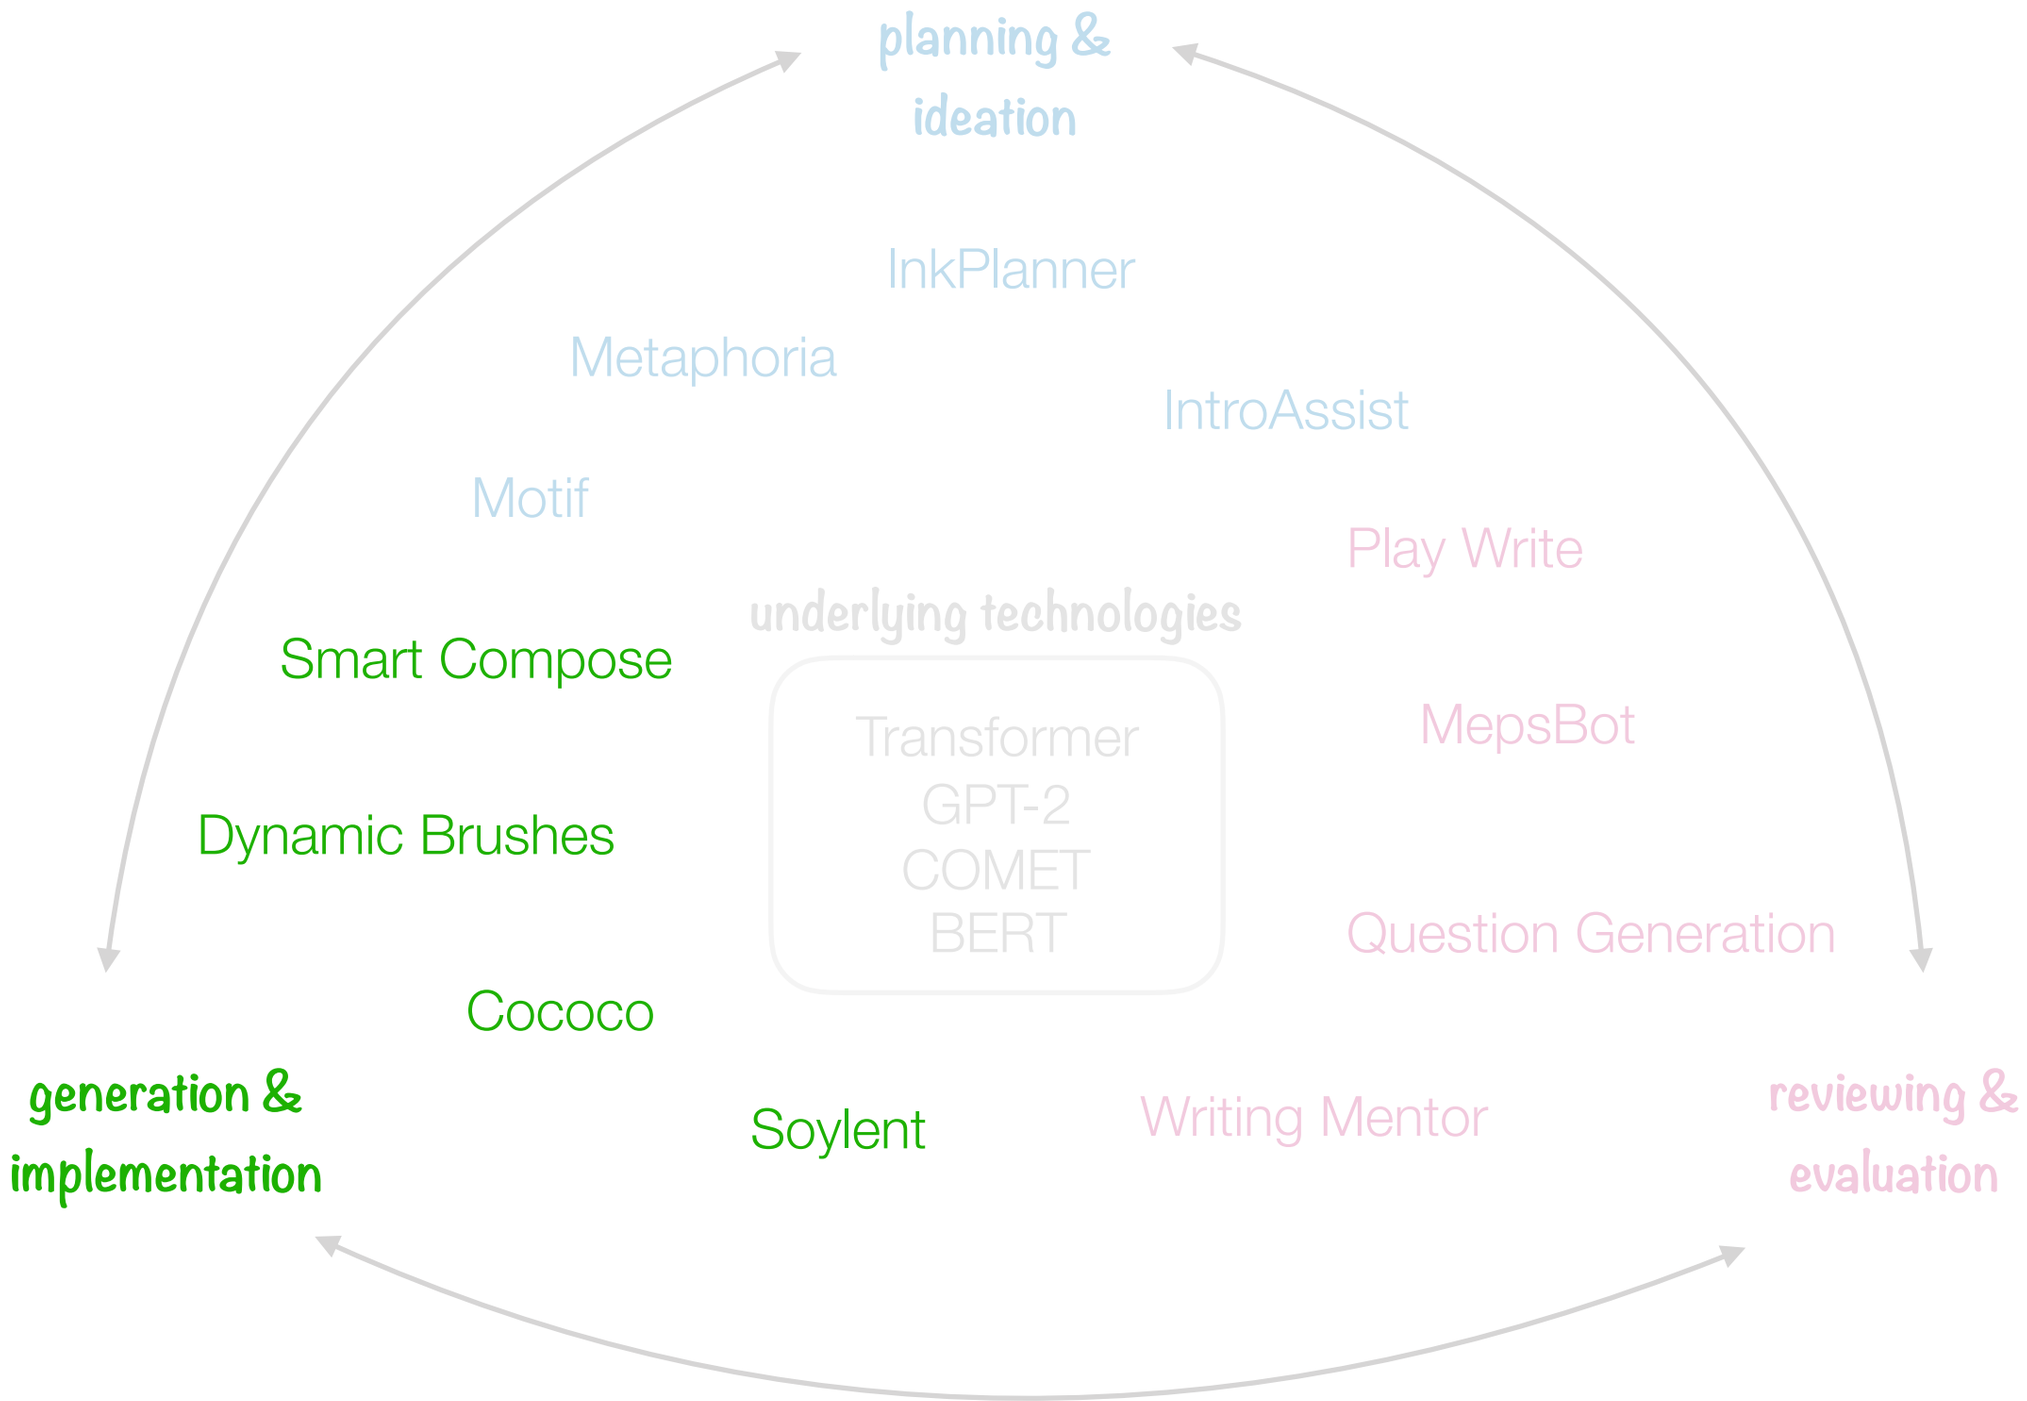

2. The “process” view of writing support.

Reviewing has seen perhaps the most work, but we still don’t know how accurate feedback has to be, or how to best present it to the writer.

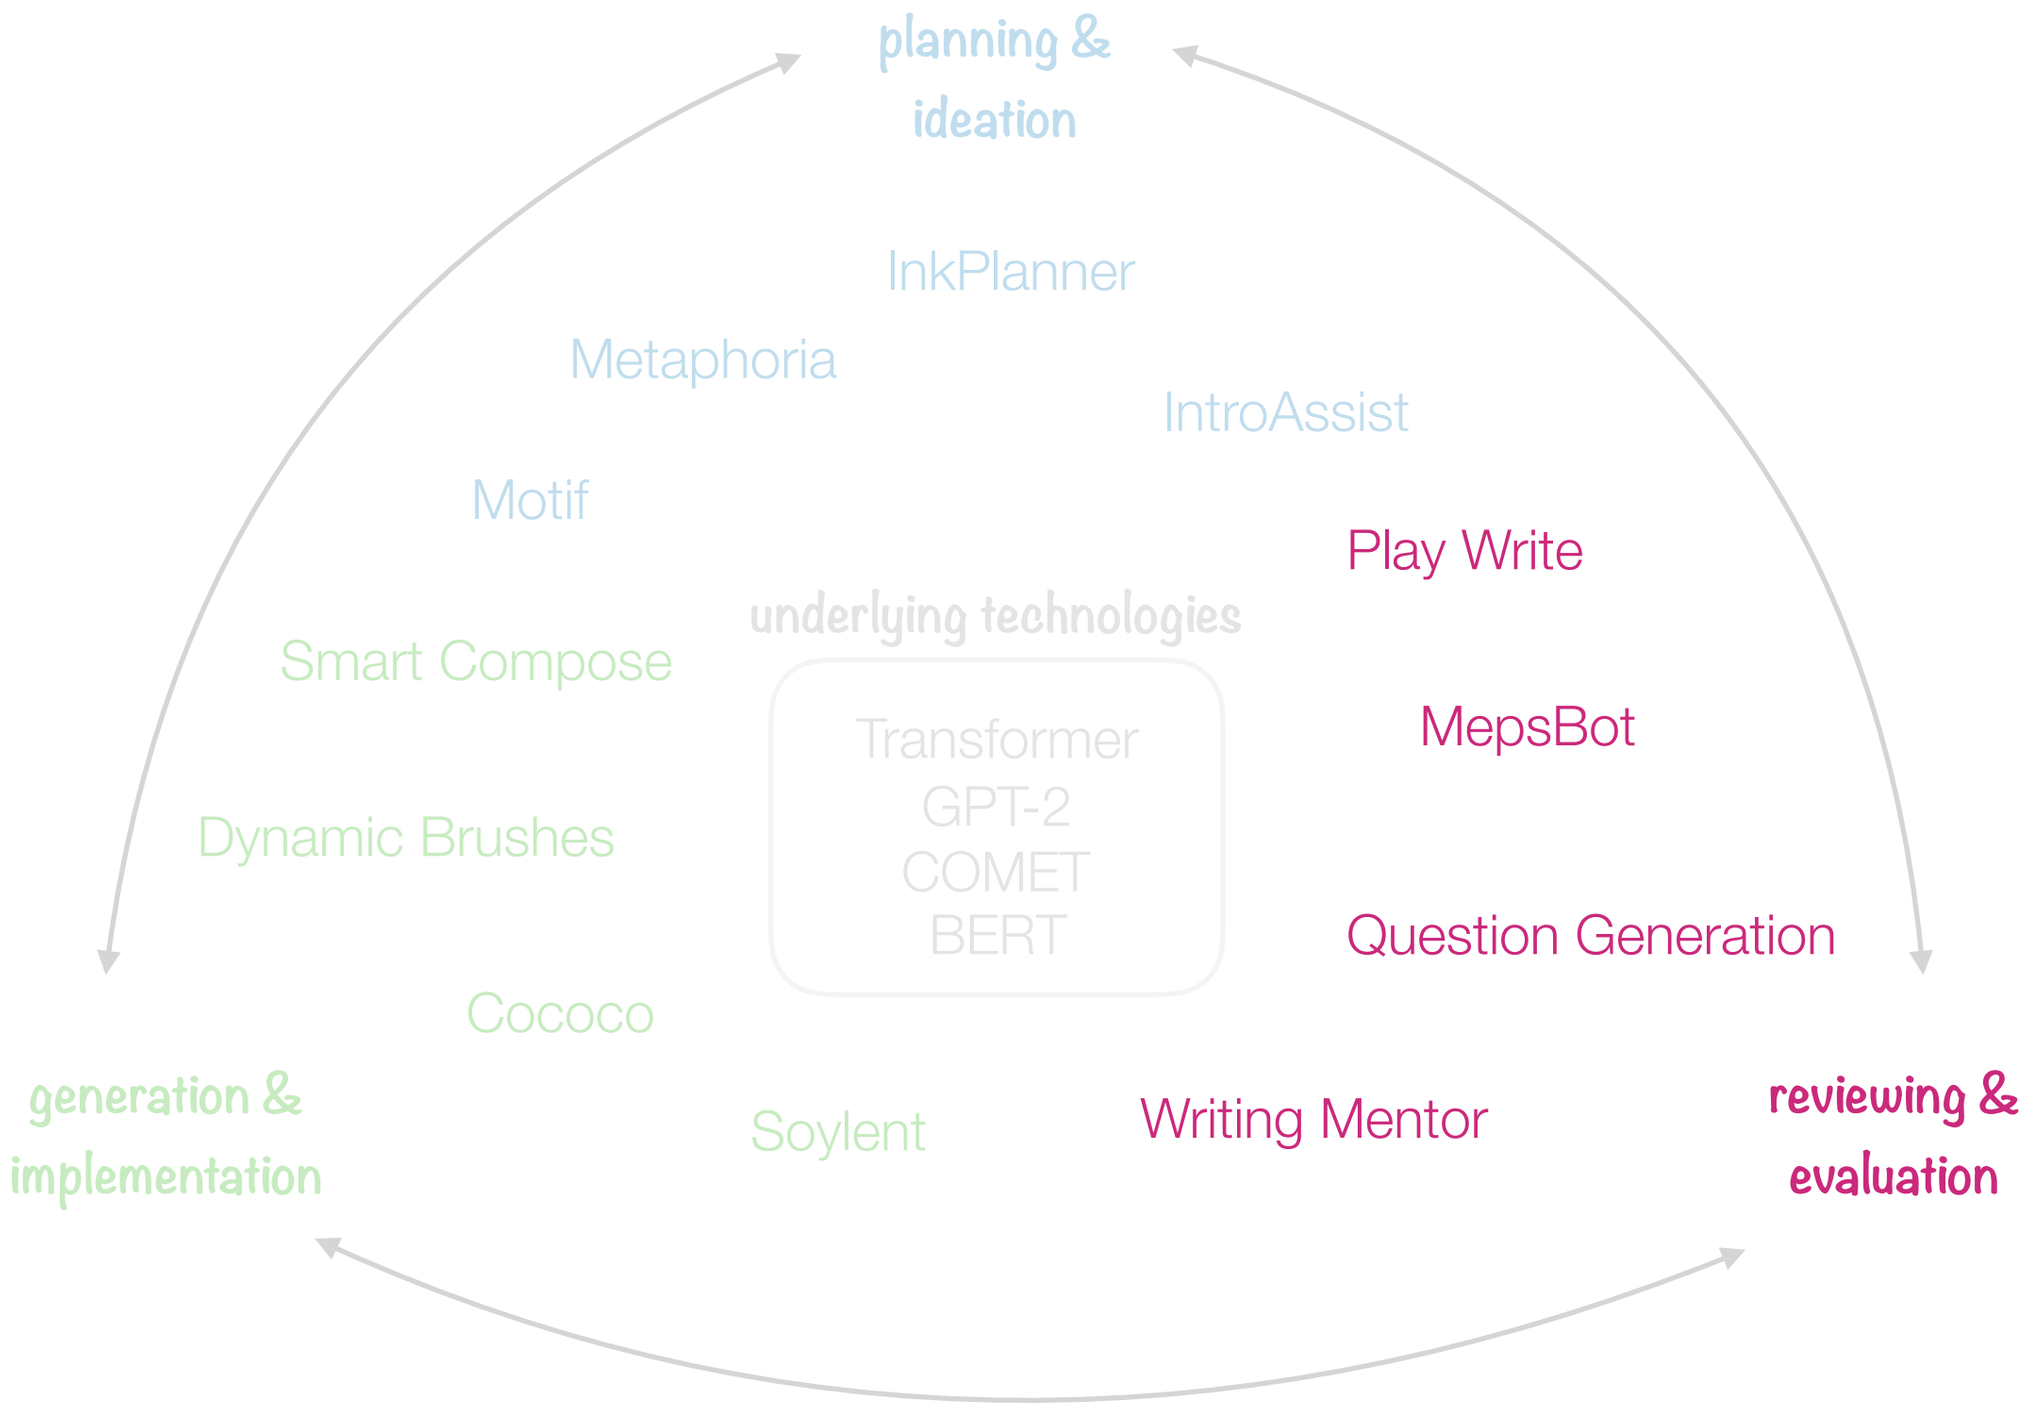

Outline

1. Theoretical underpinnings.

   Writing can be viewed as creative design, and we can define nonlinear, yet distinct, cognitive processes: planning, generation, and review.

2. The “process” view of writing support.

   Planning can most benefit from expert structures, and generation and even reviewing can benefit from generative models.

3. The “evaluation” view of writing support.

   Review of system papers from an evaluation perspective.

4. Where do we go from here?

   Reflection on where there are opportunities for new work.

Outline

1. Theoretical underpinnings.

   Writing can be viewed as creative design, and we can define nonlinear, yet distinct, cognitive processes: planning, generation, and review.

2. The “process” view of writing support.

   Planning can most benefit from expert structures, and generation and even reviewing can benefit from generative models.

3. The “evaluation” view of writing support.

   Review of system papers from an evaluation perspective.

4. Where do we go from here?

   Reflection on where there are opportunities for new work.

Design principles for creativity support tools.

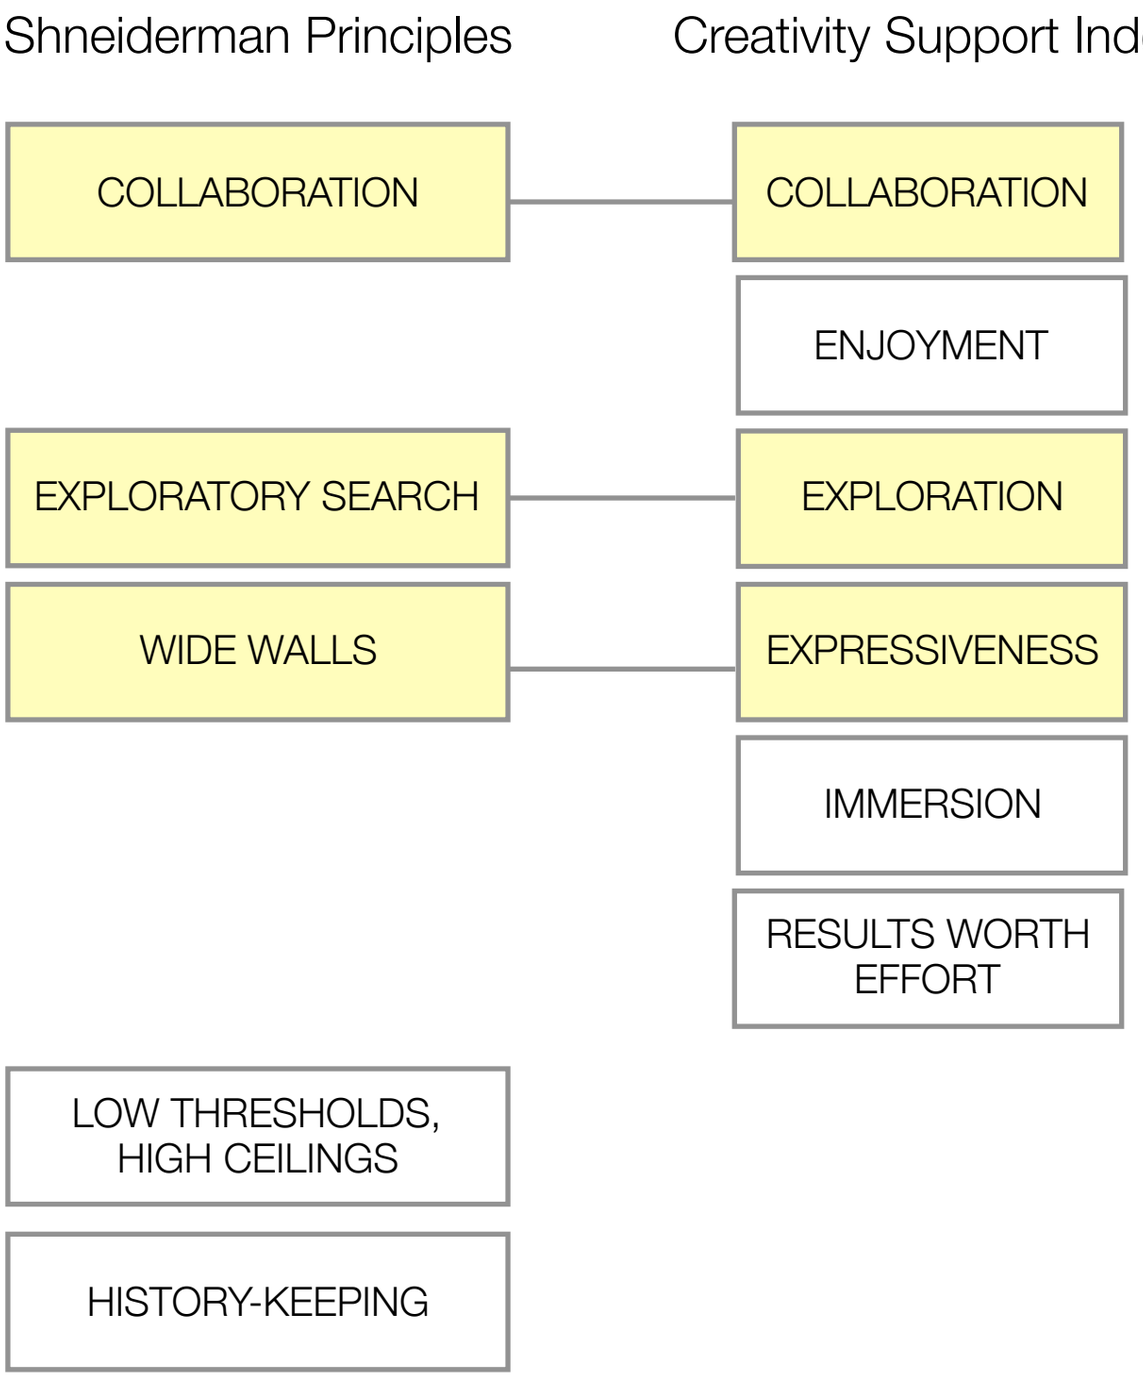

Shneiderman, 2007

Design principles for creativity support.

- inspired by work on theories of creativity:
  - enable collaboration
  - support exploratory search
  - low thresholds, high ceilings, and wide walls
  - provide rich history-keeping
- evaluation moves toward case studies

Cherry & Latulipe, 2014

A creativity support index for evaluation.

- measures and weights six factors:
  - collaboration, enjoyment, exploration, expressiveness, immersion & results worth effort
- an individual score represents the intersection of task, tool, and user

Design principles for creativity support tools.

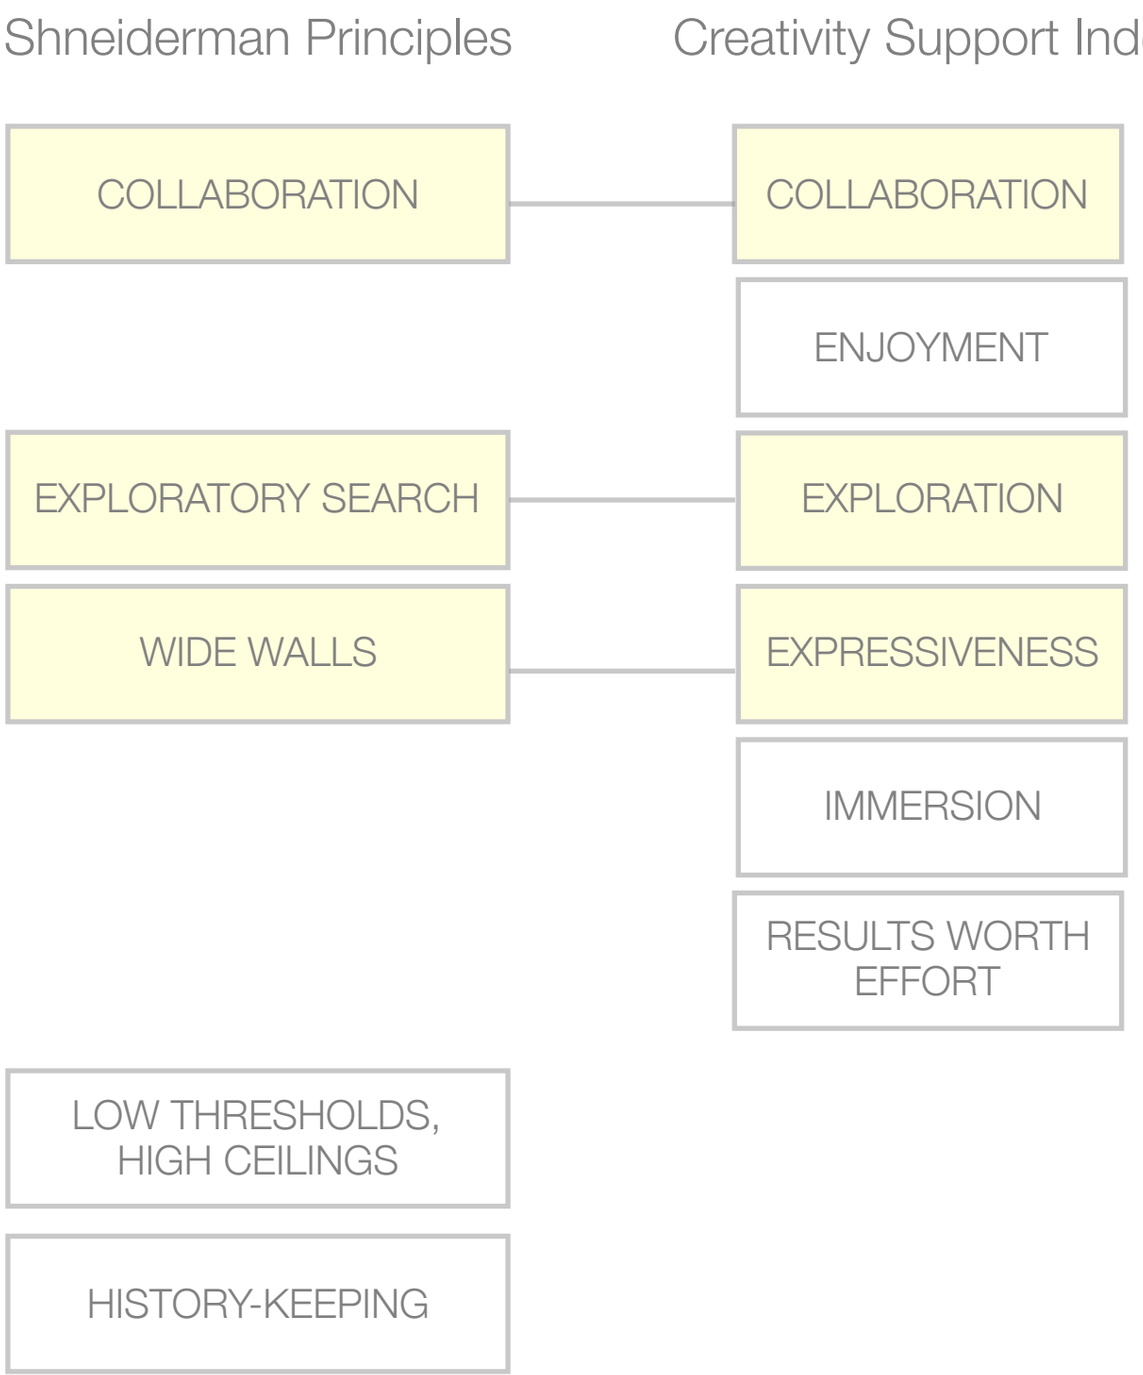

Shneiderman, 2007

Design principles for creativity support.

- inspired by work on theories of creativity:
  - enable collaboration
  - support exploratory search
  - low thresholds, high ceilings, and wide walls
  - provide rich history-keeping
- evaluation moves toward case studies

Cherry & Latulipe, 2014

A creativity support index for evaluation.

- measures and weights six factors:
  - collaboration, enjoyment, exploration, expressiveness, immersion & results worth effort
- an individual score represents the intersection of task, tool, and user

The “evaluation” view of writing support

Theory provides what tools should enable, but how people differ on how we should be assessing tools. In this section I’ll answer:

1. What are we measuring?
2. How long do we look for effects?
3. What is the comparison?

The “evaluation” view of writing support

Self-Report (9/12)	Behavior (4/12)	Product (7/12)
InkPlanner	InkPlanner	InkPlanner
Metaphoria	**Metaphoria**	**Metaphoria**
Motif	Motif	**Motif**
IntroAssist	IntroAssist	**IntroAssist**
Smart Compose	Smart Compose	**Smart Compose**
Dyn. Brushes	Dyn. Brushes	Dyn. Brushes
Cococo	**Cococo**	Cococo
Soylent	Soylent	**Soylent**
Play Write	Play Write	**Play Write**
MepsBot	**MepsBot**	**MepsBot**
Question Gen.	Question Gen.	**Question Gen.**
Writing Mentor	**Writing Mentor**	Writing Mentor

What are we measuring?

- Self-Report: how the creator feels
  - how hard was the task? (NASA-TLX)
  - do they like the tool?
  - did they feel ownership of the product?
- Behavior: how the creator acts
  - do they use suggested content (edit distance)
  - patterns of feature usage
- Product: what the creator makes
  - how much they make
  - expert ratings of final product

We need more focus on how the creator acts, which can validate their self-report and guide us towards why the final product improves.

The “evaluation” view of writing support

Short term	Long term
InkPlanner	InkPlanner
Metaphoria	Metaphoria
Motif	Motif
IntroAssist	IntroAssist
Smart Compose	Smart Compose
Dyn. Brushes	**Dyn. Brushes**
Cococo	Cococo
Soylent	Soylent
Play Write	Play Write
MepsBot	MepsBot
Question Gen.	Question Gen.
Writing Mentor	**Writing Mentor**

How long do we look for effects?

- one time (about one hour) studies
  - introassist task can be done quickly, and people may find the tool less useful over time.
- long-term (repeated visits) studies
  - dynamic brushes is a complex tool that takes time to learn and integrate into an artistic practice.

We need more focus on how tools impact use outside the lab and with long-term usage.

The “evaluation” view of writing support

No comp.	Ablation	Existing	No access
InkPlanner	InkPlanner	InkPlanner	InkPlanner
Metaphoria	Metaphoria	Metaphoria	**Metaphoria**
Motif	**Motif**	Motif	Motif
IntroAssist	IntroAssist	IntroAssist	**IntroAssist**
Smart Compose	Smart Compose	Smart Compose	Smart Compose
Dyn. Brushes	Dyn. Brushes	Dyn. Brushes	Dyn. Brushes
Cococo	**Cococo**	Cococo	Cococo
Soylent	Soylent	Soylent	Soylent
Play Write	Play Write	**Play Write**	Play Write
MepsBot	MepsBot	MepsBot	**MepsBot**
Question Gen.	Question Gen.	Question Gen.	Question Gen.
Writing Mentor	Writing Mentor	Writing Mentor	Writing Mentor

What is the comparison?

- no comparison or gold standard result
  - dynamic brushes has no comparative tool
- ablated tool
  - cococo removes the steering tools
- existing tool
  - play write compares to standard microsoft word
- no access to tool
  - mepsbot compares pre- & post-feedback writing

Comparison is very dependent on what the tool provides; since many tools tackle niche creative acts, often there are no existing tools to compare against.

Outline

1. Theoretical underpinnings.

   Writing can be viewed as creative design, and we can define nonlinear, yet distinct, cognitive processes: planning, generation, and review.

2. The “process” view of writing support.

   Planning can most benefit from expert structures, and generation and even reviewing can benefit from generative models.

3. The “evaluation” view of writing support.

   In current work there is lack of behavioral and long-term evaluations.

4. Where do we go from here?

   Reflection on where there are opportunities for new work.

Outline

1. Theoretical underpinnings.

Writing can be viewed as creative design, and we can define nonlinear, yet distinct, cognitive processes: planning, generation, and review.

2. The “process” view of writing support.

Planning can most benefit from expert structures, and generation and even reviewing can benefit from generative models.

3. The “evaluation” view of writing support.

In current work there is lack of behavioral and long-term evaluations.

4. Where do we go from here?

Reflection on where there are opportunities for new work.

Outline

1. Theoretical underpinnings.

   Writing can be viewed as creative design, and we can define nonlinear, yet distinct, cognitive processes: planning, generation, and review.

2. The “process” view of writing support.

   Planning can most benefit from expert structures, and generation and even reviewing can benefit from generative models.

3. The “evaluation” view of writing support.

   In current work there is lack of behavioral and long-term evaluations.

4. Where do we go from here?

   Reflection on where there are opportunities for new work.

What is works and what is lacking for each process?

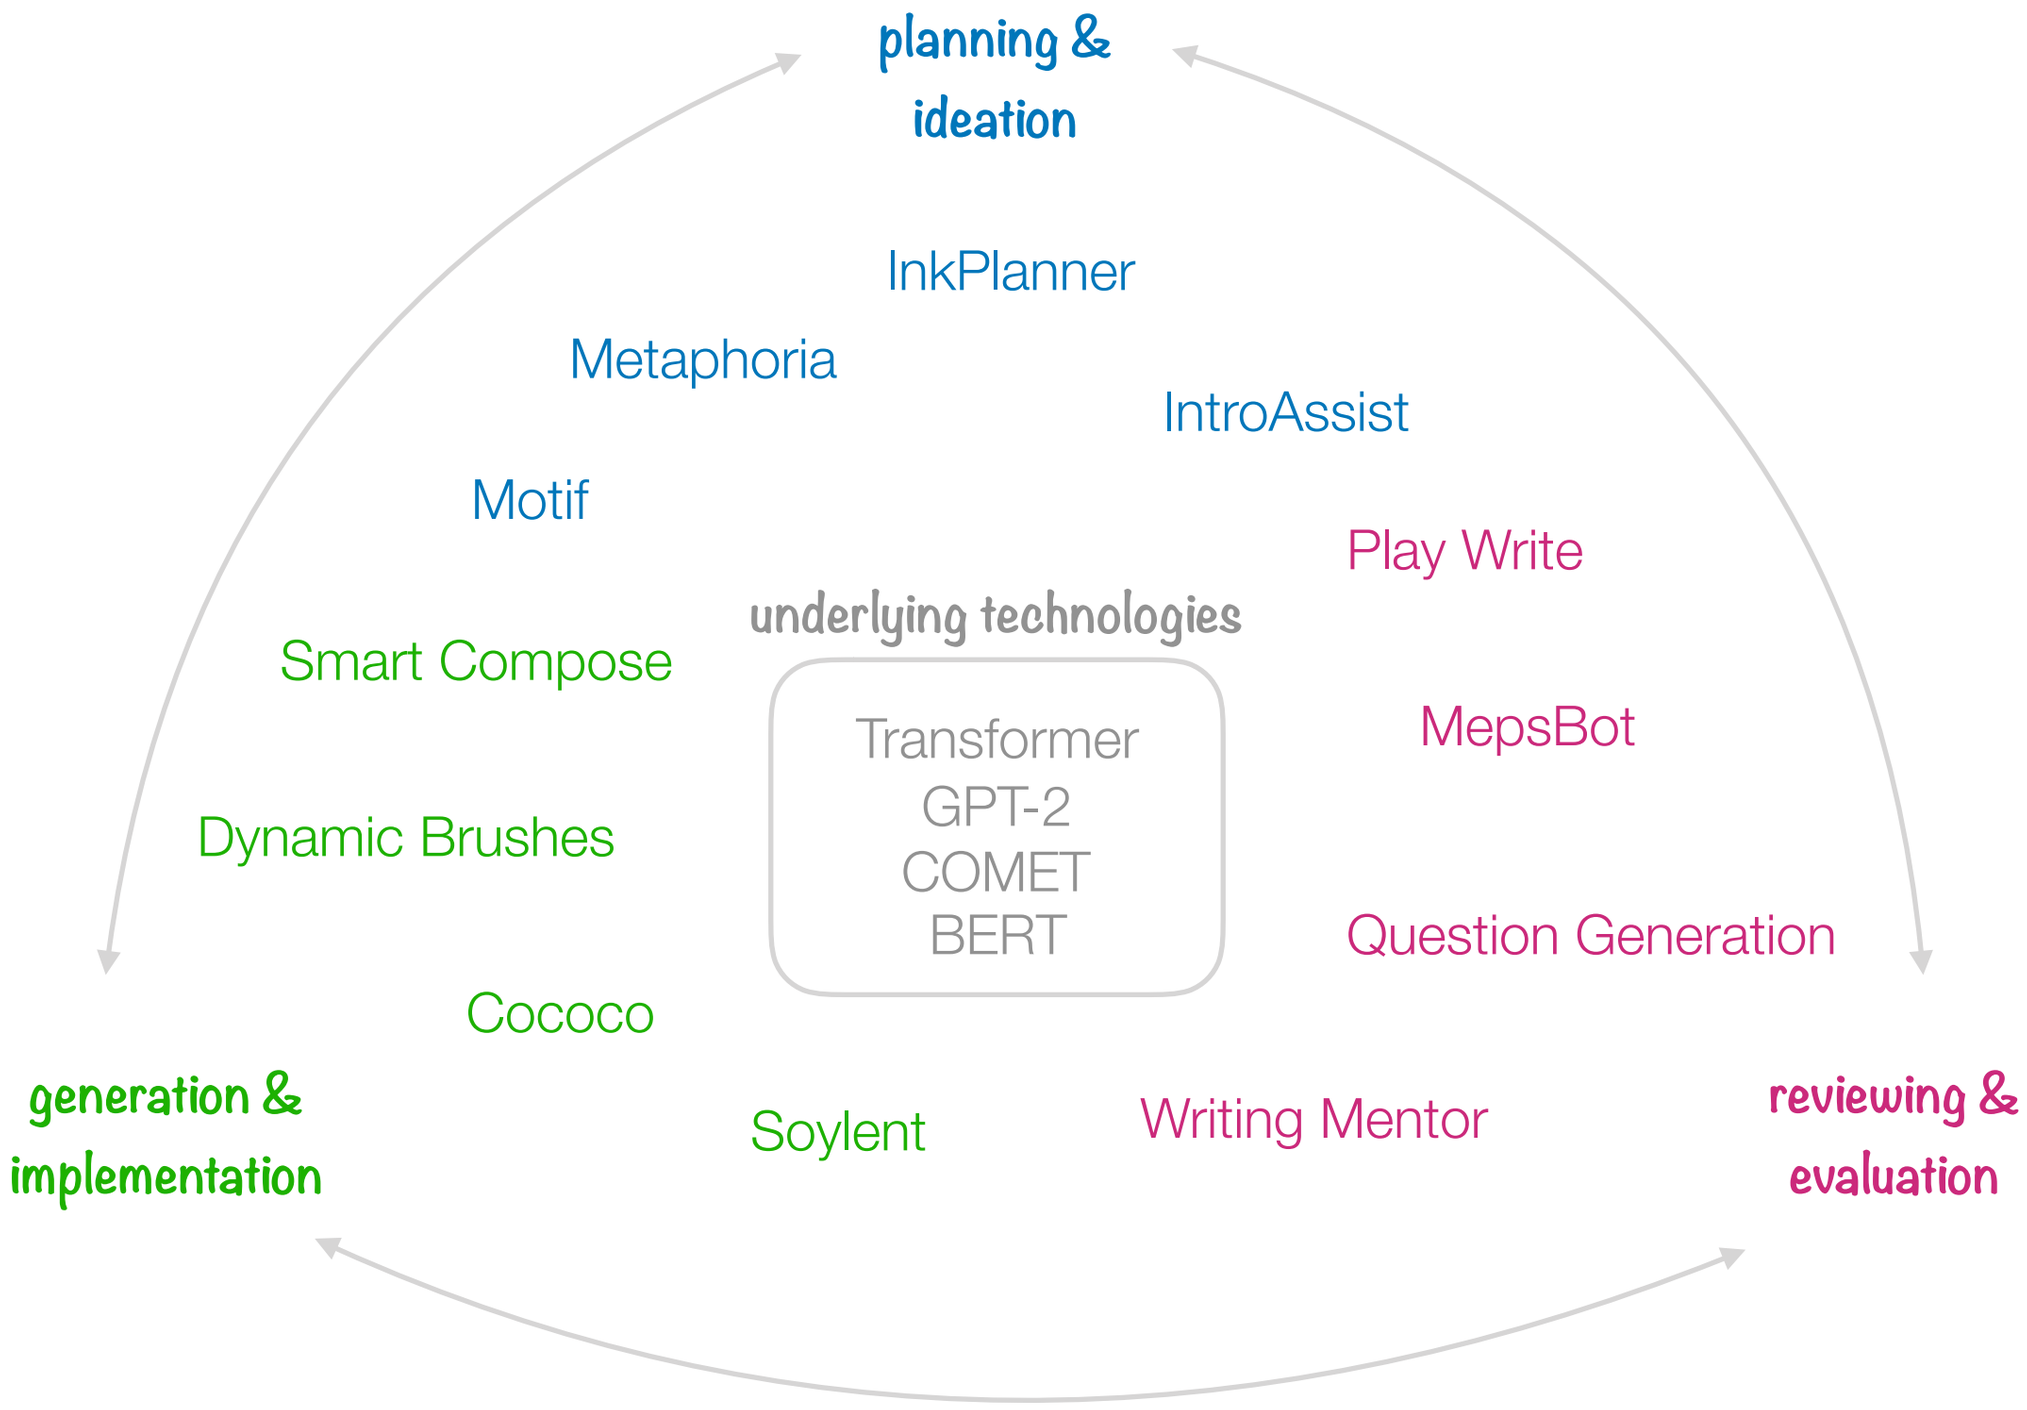

What is works and what is lacking for each process?

planning &
ideation

generation &
implementation

reviewing &
evaluation

What is works and what is lacking for each process?

planning &
ideation

Right now, planning benefits most from expert structures.
Ideation tools often fail to be semantically coherent.

Generative tools have promise but require new interfaces.

generation &
implementation

Reviewing has seen perhaps the most work, but many questions remain.

reviewing &
evaluation

Expert structures will best support planning.

- Writers rely on genre conventions when planning, but these expert structures are difficult to distill, learn and apply.
  - distill: How can we extract expert structures from unstructured text?
  - learn: How do we best teach these structures to novices?
  - apply: How do we help writers apply structures to their own work?
- Making these structures explicit, and guiding their usage, lowers the threshold for novices.
  - Can these structures help collaboration?
  - Can they help experts, who may already implicitly know them?

Language models will support generation & reviewing.

- Language models have clear potential to support generation. Iteration requires making small changes based on what a writer has learned from a draft.
  - How can we give writers iterative control over text generation?
- They can be used for reviewing specific writing qualities. Language models are often the back-bone of text classifiers, which can be used to review specific qualities of text, like the emotional content or the concreteness.
  - How well can classifiers predict these qualities?
  - How accurate do these models need to be to provide useful feedback to writers?
  - What's the best way to present this feedback to writers?

What is the state of HCI work on writing support?

1. Theoretical underpinnings.

Writing can be viewed as creative design, and we can define nonlinear, yet distinct, cognitive processes: planning, generation, and review.

2. The “process” view of writing support.

Planning can most benefit from expert structures, and generation and even reviewing can benefit from generative models.

3. The “evaluation” view of writing support.

In current work there is lack of behavioral and long-term evaluations.

4. Where do we go from here?

Expert structures and generative language models as writing support.

What is the state of HCI work on writing support?

1. Theoretical underpinnings.

Writing can be viewed as creative design, and we can define nonlinear, yet distinct, cognitive processes: planning, generation, and review.

2. The “process” view of writing support.

Planning can most benefit from expert structures, and generation and even reviewing can benefit from generative models.

3. The “evaluation” view of writing support.

In current work there is lack of behavioral and long-term evaluations.

4. Where do we go from here?

Expert structures and generative language models as writing support.

🎊 ✨ *Thank you for listening.* 💫 🎉

Fin.